Everything For Electronics

Nuts & Volts

October 2005
www.nutsvolts.com

WAXING DIGITAL WITH ELECTRONIC WICKS

Prop-1 Controller
PARALLAX EFX

I Love My Heils!

Wireless Mesh Networks

Build an MP3 Player

Nuts & Volts

OCTOBER 2005

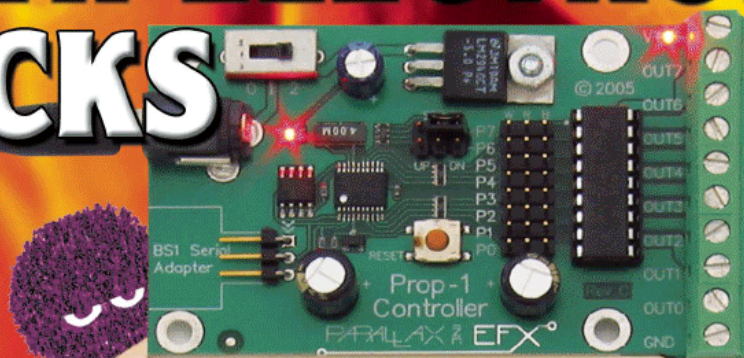Vol. 26 No. 10

p. 14

p. 67

p. 26

p. 46

ROBOTICS MICROCONTROLLERS COMPUTER CONTROL SPACE
Everything For Electronics
Nuts & Volts
October 2005
www.nutsvolts.com
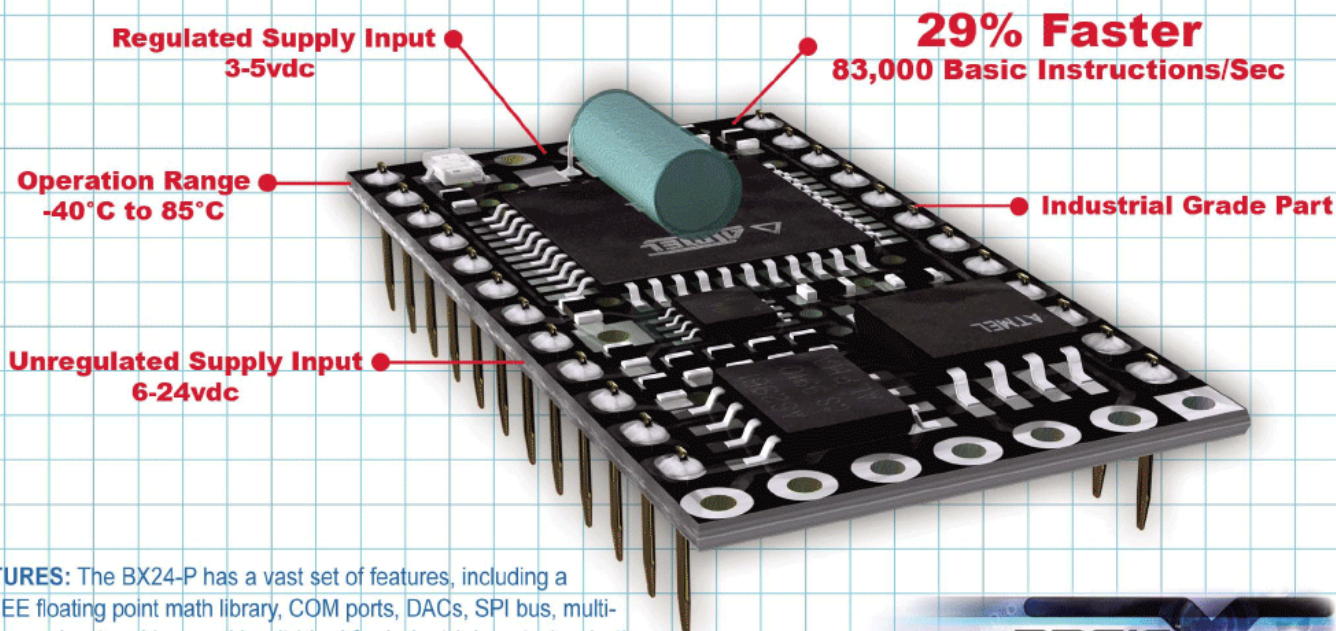WAXING DIGITAL WITH ELECTRONIC WICKS
I Love My Heils!
Wireless Mesh Networks
Build an MP3 Player

Link Instruments

PC-Based Test Equipment

Logic Analyzers

New LA-5000 Series

- 40 to 160 channels
- up to 500 MSa/s
- Variable Threshold
- 8 External Clocks
- 16 Level Triggering
- up to 512K samples/ch
- USB 2.0 and Parallel Interface
- Pattern Generator option

LA5240 (200MHz, 40CH)	$1700	USB 2.0/Parallel
LA5280 (200MHz, 80CH)	$2350	USB 2.0/Parallel
LA5540 (500MHz, 40CH)	$2500	USB 2.0/Parallel
LA5580 (500MHz, 80CH)	$3500	USB 2.0/Parallel
LA55160 (500MHz, 160CH)	$7500	USB 2.0/Parallel

Small and portable LA-2124

$800

- Small, Lightweight and Portable
- Only 4 oz and 4.75" x 2.75" x 1"
- Parallel Port Interface to PC
- Trigger Out

- 24 Channel Logic Analyzer
- 100MSa/S max sample rate
- Variable Threshold Voltage
- Large 128k Buffer

Digital Oscilloscopes

- 2 Channel Digital Oscilloscope
- **100 MSa/s** max single shot rate
- 32K samples per channel
- Advanced Triggering
- Only 9 oz and 6.3" x 3.75" x 1.25"

• Small, Lightweight, and Portable	DSO-2102S	$525
• **USB or Parallel Port** interface	DSO-2102M	$650
• Advanced Math	DSO-2102S(USB)	$600
• FFT Spectrum Analyzer (option)	DSO-2102M(USB)	$725

Link Instruments (973) 808-8990

17A Daniel Road East · Fairfield, NJ 07004 · Fax (973) 808-8786

www.Link-instruments.com

Reader Feedback

Dear Nuts & Volts:

In the August issue, page 15: I wanted to ask how the author, was able to complete the project as he states "it only cost me about $50.00 ..." On page 50 of the July issue, I looked up the parts he listed on the kronosrobotics.com website and for just their parts (not counting other required suppliers' parts), the total came to at least $133.00. Please explain this quote.

Charles E. Morris

Author response:

Charles, if you take a closer look at the UUM article you will notice that the $50 cost is referring to the cost of the first analyzer project I completed. The original analyzer used the PC for display. The PC software is available free from the Kronos Robotics website at **www.kronosrobotics.com**

If you build your own board, you can actually build that analyzer for less than $32. With a carrier board kit and Dios chip you can build it for $51.95. The UUM project will cost you over $150 if you have to purchase all components and hardware. However, depending on the hardware on hand and the type of enclosure, it could be built for less than $100.

Michael Simpson

Dear Nuts & Volts:

I wish to inform you your August cover story has a spelling error. The correct terminology for phone hacking is phreaking not phracking. The person would then be a phreaker not a phracker and to say your data was phracked or phreaked does not make much sense. I have been in the security field for a while, but to be fair, I thought there may be some regional differences in spelling but I could find no references to it spelled phracking.

Justin Derleth

Author response:

Phrack was originally the name of a long-running, popular magazine for Phreakers and Hackers, which coincidentally closed up shop with its last issue just recently. Two infamous hackers started the magazine, combining the words Phreak and Hack to create the magazine's title, Phrack (see **http://en.wikipedia.org/wiki/Phrack** *Today, due to the popularity of the magazine, the hacker underground and others occasionally use the words Phrack, Phracker, Phracking, and Phracked (search for each individually using Google) as bastardized terms to mean anything from hacking in general to phone cracking or phreaking in particular.*

Phreak is probably more accurate. Given the close ties between the words Phrack and Phreak, I probably had a Freudian Slip or whatever you want to call it and wrote Phrack as the root word instead of Phreak, which is the correct formal term for phone cracking.

David Geer

An error occurred in the "Home-Brew Power Supply" schematic on page 13 of the Aug. 2005 issue in the Q&A. Here is the corrected schematic. Sorry for any inconvenience. The original schematic will work, but it works better with the correction. — TJ

Commercial Power Supply

$$R2 = \frac{240Vout - 300}{1.25}$$

Home-Brew Power Supply

TechKnowledgey 2005

Events, Advances, and News From the Electronics World

If you see anything revolutionary and interesting — or just plain cool —while surfing newsgroups, filtering press releases, or simply doing your job, drop me a line about it. Have questions or comments about what you read here in "Techknowledgey 2005?" Send those my way, too. You can reach me at **www.jkeckert.com**

— *Jeff Eckert*

Advanced Technologies
New Tool for Geoanalysis

The left half shows the simulated geoneutrino production distribution detectable by KamLAND, and the right half shows the Earth's structure. Photo courtesy of Stanford University.

In the past, if you were inclined to wonder about the internal workings of the Earth, seismology (the analysis of vibrations produced by earthquakes and sensed by thousands of instrument stations worldwide) was pretty much the only analysis tool available to you. But a new approach has been developed by physicists who study neutrinos and their antiparticles, called "antineutrinos," which emanate from nuclear reactors.

As it turns out, antineutrinos also flow from the Earth's interior when uranium and thorium isotopes undergo a cascade of heat-generating radioactive decay processes. A detector in Japan, called KamLAND (for Kamioka liquid scintillator antineutrino detector) has sensed the geologically produced antineutrinos, known as "geoneutrinos."

Lying in a cavern beneath a mountain that shields the installation from the background noise of cosmic radiation, KamLAND consists of about 2,000 photomultiplier tubes, each 20" (51 cm) in diameter and contained in a 59-ft (18-m) vessel, bathed in 1,000 tons of liquid scintillator. It was actually built back in 1997 for more academic purposes.

This new application for KamLAND, developed in collaboration with physicists from Stanford University (**www.stanford.edu**) and the Lawrence Berkeley National Laboratory (**www.lbl.gov**), could eventually yield important geophysical information. Seismology provides data about the locations of boundaries of different types of rock, but not much else. Geoneutrinos, in contrast, provide some information about how much uranium and how much thorium lies below. It doesn't reveal anything about the crystal structure, e.g., whether the thorium is thorium oxide or thorium nitride. But with further development, the process could still be highly useful. Because radioactive heat drives plate tectonics, getting accurate ratios of thorium to uranium isotopes will help scientists better understand deep-Earth processes.

But don't expect any major headlines in the short term. According to Stanford physics professor Giorgio Gratta, "It's a revolution, but let me temper this a little bit with the physi-cists' point of view. Those are very difficult measurements, and those detectors are very expensive and large. So before the revolution really comes to fruition, I think it'll take some time, I would imagine one or two decades, before we have more of those detectors and maybe larger ones built in the appropriate place for geophysics."

Laser Device Senses Lethal Chemicals

Prof. Boris Mizaikoff displays a prototype of the gas phase sensor, while graduate student Christy Charlton holds the liquid phase prototype. Photo courtesy of Georgia Tech.

Advancements in liquid and gas phase chemical sensing being made at the Georgia Institute of Technology (**www.gatech.edu**) may lead to the development of palm-sized sensing tools that can provide instant detection to foil attempts by terrorists to lace water supplies with lethal chemicals, for example. In a large city, tiny liquid phase sensors located at strategic points in the city's water mains could detect a chemical as it passes and tell a computer to close

down the affected pipes.

Using small quantum cascade lasers, researchers at Georgia Tech, with colleagues from Tel-Aviv University (**www.tau.ac.il**) and OmniGuide Communications (**www.omni-guide.com**), have built and demonstrated a prototype handheld gas phase chemical sensing device, as well as a liquid phase sensing device.

The quantum cascade laser is the key component. Almost every organic molecule has a distinctive absorption pattern in the mid-infrared range (roughly between 3 and 20 microns). Illuminating these molecules with a laser that is tuned to their fingerprint frequency will cause them to vibrate as they absorb radiation at that frequency.

Detecting a chemical therefore becomes as simple as illuminating a small volume of gas or liquid with a laser. If the laser is tuned to a characteristic absorption frequency of benzene, for example, and benzene is present, the molecules will vibrate and absorb an amount of radiation at its characteristic absorption frequency, indicating its concentration.

For the gas-sensing modules, the researchers use a photonic band gap hollow waveguide (developed by OmniGuide), which is essentially a hollow, flexible tube to both contain very small amounts of the air being sampled and assist in sensing. The waveguide can be built to propagate only one wavelength of light. Therefore, when the laser illuminates the gas molecules inside the waveguide, the waveguide will propagate only the selected fingerprint frequency for detecting a specific molecule. Further improvements in sensitivity are expected to result in instruments that provide parts-per-billion precision.

Computers and Networking
New Dual-Core Processor Introduced

With an eye toward the Christmas season, Advanced Micro Devices (AMD, **www.amd.com**) has introduced the Athlon™ 64 X2 dual-core processor 3800+, which will appear in PCs in time for holiday buying. Priced at $354 in manufacturing quantities, it is designed for entry-level users who need multi-tasking capabilities and high performance on digital media, exceeding what can be offered by single-core processors running multiple applications simultane-

ously. Higher-end versions are priced from $537 to $1,001, depending on the model and performance.

In the AMD announcement, no clock rate was mentioned, but the company does note correctly that, "Frequency alone is not an accurate measure of performance." The "+" in the model number indicates added performance benefits delivered by the architecture, such as an integrated memory controller and HyperTransport technology, but you will just have to wait for the computer magazines to come up with the results of benchmark tests for a comparison to other devices.

PCs Drop Below $200

It was bound to happen eventually: a desktop PC for less than $200. As of this writing, Tiger Direct (**www.tigerdirect.com**) is offering a new system for $179. For this paltry price, you get a Chaintech Barebone motherboard with a 2167 MHz AMD Athlon 3000+ processor installed in a Soyo mid-tower case with a 400 W power supply. It also includes 512 MB of RAM, a cooling fan, and a keyboard and mouse. The price includes a $90 rebate that is set to expire before this issue goes into print, so the deal may not be available by the time you read this. But look for competitors to join in the continuing downward price spiral.

Circuits and Devices
Spherical Semiconductors?

Over the years, I have received press releases about some pretty nutty things, including a motor that was purported to be capable of powering an automobile using four D cells and a hat that, when worn by your dog, would allow him to converse with you in English. (If only someone could develop such a hat for New York cab drivers.) At first glance, a concept being promoted by Ball Semiconductor, Inc. (**www.ballsemi.com**), might appear to fall into the same category, and yet it sounds plausible for some limited applications.

The concept is to eliminate con-

TechKnowledgey 2005

temporary semiconductor technology, which consists of a complex and expensive series of steps from crystal formation, wafer fab, photolithography, and assembly, and replace it with a process that produces single semiconductors in a continuous process. The kicker is that such a device would be in the

Concept of a 1-mm spherical semiconductor. Photo courtesy of Ball Semiconductor, Inc.

form of a 1 mm sphere, rather than a flat, rectangular chip. Using a manufacturing line that consists of small tubes and pipes, the devices would be generated in a continuous process using gases and chemical reactions.

According to the company, "Our process begins with the sorting of very small polycrystal silicon granules that are then processed into single-crystal silicon balls. Initially, the company is producing 1 mm single-crystal balls; but we also intend to develop the capability to produce smaller balls. We anticipate that the balls will be in constant motion as they are processed, treated, and transported at high speed through hermetically sealed pipes and tubes during various processes for crystal-growing, grinding and polishing steps, as well as for the repeated cleaning, drying, diffusion, film deposition, wet and dry etching, coating, and exposing steps of the integrated-circuit manufacturing process. The spheres are exposed to air only during photolithography; thus, there is no need for the traditional — and expensive — clean room."

This brings up questions about how one might attach leads to the devices, mount them on a motherboard, or provide power to them, but they might be capable of functioning as self-contained, disposable units. Indeed, the company envisions a prime application in the form of "MedBalls," which would function as ingested or implanted devices for medical diagnosis and treatment. Is this for real? Time will tell.

Clean Backup Battery Developed

At present, if you want to keep your PC or other equipment running through a power outage, you'll need to plug it into a relatively bulky and expensive uninterruptible power supply. But NEC Corp. (**www.nec.com**) has come up with an alternative approach that it calls the organic radical battery (ORB). The ORB — consisting of four cells — is capable of driving a 140 W desktop computer. It can't keep the equipment running for more than a few tens of seconds, but that is usually sufficient for saving open files and executing a normal shutdown. And because the ORB is only 55 x 43 x 4 mm in size and weighs only 88 g, it can be built into the equipment.

OCTOBER 2005

Unlike other rechargeable batteries, the device is environmentally friendly, being based on an organic polymer that contains no dangerous heavy metals, and it is said to be non-flammable and non-explosive. The ORB is still in the development stage and not yet commercially available, but watch for it to appear in systems where protection against data loss is essential.

Industry and the Profession
Spammer Pays Up

In wrapping up a 2003 lawsuit filed by Microsoft Corp. (**www.micro soft.com**) against "spam king" Scott Richter and his company, OptIn RealBig.com LLC, Richter has agreed to pay $7 million to Microsoft. The settlement also stipulates that Richter, his company, and his affiliates will continue to comply fully with all federal and state anti-spam laws, including the US CAN-SPAM (Controlling the Assault of Non-Solicited Pornography and Marketing) Act, and will not send spam to any person who has not confirmed a willingness to receive the email.

Reportedly, at his peak, Richter was responsible for the distribution of more than 38 billion pieces of unsolicited commercial email per year. (To put it into perspective, that is equal to 128 spams per person in the US.) Microsoft will direct $5 million of the settlement to expand the company's Internet safety partnerships with governments and law enforcement agencies worldwide through technical training, investigative and forensic assistance, and the development of new technology tools.

The company has pledged an additional $1 million to provide many community centers in New York State with broader access to computers for underprivileged children and adults through Microsoft's Unlimited Potential Program.

Dell to Open New Customer Contact Center

If you are tired of calling customer support numbers and enduring communication problems with people whose English is buried under a heavy Hindi/Urdu accent, you may be in for a new experience. Starting next year, when you call for assistance with your Dell PC, you will be able to endure communication problems with people whose English is buried under a heavy Filipino accent.

Dell has announced that it will open a customer contact center in the Metro-Manila area early next year to provide technical and customer support for consumer customers. Initial plans call for the hiring of approximately 700 technical support and customer service agents. **NV**

Putting the Spotlight on BASIC Stamp Projects, Hints, and Tips

Stamp Applications

Tricks and Treats With LEDs

October — it's finally here! It's time to build (okay, finish building) our Halloween displays — whether they're used in yard decorations or in a full-blown, professional haunt. Nothing adds spooky ambience like a candle in a darkened room, but candles can be dangerous — unless you build them yourself and substitute the coolness of LEDs for the heat of an open flame.

A few months ago my friend Brian Bayliss — who runs the Methodz of Madness Halloween prop-building forum — asked me if I could help him create a set of faux candles using a Prop-1 (BS1) controller. "Sure!" I quickly exclaimed, and in short order I emailed a bit of code for him to test. After a few rounds of fine-tuning, Brian was pretty happy.

But it didn't look quite right. Nobody has ever accused the BS1 of being a speed demon, so I had to keep things simple. And that's exactly what I did: I tumbled a random number and spit it out to the pins (that were connected to LEDs). While it looked *pretty good*, there was something distinctly "digital" about it. And no wonder — all I was doing was turning LEDs on and off.

Then my colleague John Barrowman stepped in and added a capacitor to each "wick" output. Bam! Now we have something. The capacitor doesn't change the turn-on time of the circuit, but it does cause the LED to fade off instead of simply snapping off. It's amazing how something so simple can have such a dramatic visual impact.

Back in July, John, Jen Jacobs (our graphic designer), and I attended the Midwest Haunters Convention in Columbus, OH. It was great time with lots of interesting people and the opportunity to show some amazingly skilled prop builders how sim-

ple controllers like the BS1 can add real life to their displays. The booth next to us had a hanging man that was animated, but not very "animated." John asked the owner if we could add a controller to the prop and he agreed. While John wired up the Prop-1, I wrote a simple program and after about 15 minutes the hanging man prop was thrashing around as if he'd actually been lynched. It was fun watching the crowd gather, and the prop owner was smiling for the rest of the show. (Interestingly, he had purchased a few Prop-1 controllers from Parallax, but hadn't learned how to use them. This demo showed him what one could do with very simple code.)

Our booth didn't have any thrashing props, but the one thing we did have was a huge hit: our faux candles. In fact, some folks were irritated that we weren't offering them for sale! Figure 1 gives you an idea of what a completed candle looks like.

The candles require very little effort to put together, and I'm going to give you all the details — mechanical and electrical. Now, if you're skeptical about how good an electronic candle can look, perhaps this will sway you: on three separate occasions we were told we could not have

Figure 1. Electronic Candle.

an open flame on the trade show floor! Our response was to put a hand over the faux candle — in a manner that one could never do with the real thing — and then ask, "What open flame?"

Waxing Digital

Okay, how do we get the electronic wicks into the candles? It's really pretty easy, and we'll do that first so that we can fit the electronics to it. Run down to your local dollar store and buy a bunch of cheap candles. We tend to go with the short, stubby kind, and find that lighter colors work best with yellow LEDs (*Note: While white LEDs seem like a good idea, we've tried them and didn't care for the results. Plus, they're really expensive, so don't bother with them*). When you get home, set up

your candles in a safe location and light them. What we want to do is allow the wicks to burn wells about 1" deep into the tops of the candles. When that's complete (it may take a couple of hours), extinguish the candles and let them cool completely.

Now, take a 1/2" drill bit and bore a hole down the center of the body. After the centerline holes are drilled, add a small hole between the wick channel and the outside edge of the body to hold the capacitor. You may also want to cut a small notch into the bottom of the candle to route wiring.

Electronic Wicks

Figure 2. Candle Schematic.

Now that we have candle bodies, let's build some electronic wicks to light them. The circuit is quite straight-forward: we drive three yellow (you could mix in an orange too) LEDs using a ULN2003 (so we don't overburden the BASIC Stamp). Figure 2 shows the circuit.

Depending on the LEDs and power supply you use, you may want to adjust the resistor value. We've found that calculating 20 mA through the LEDs works well, and three LEDs per candle body produce very nice results. Of course, if your LEDs can take more continuous current and you'd like things to be brighter, go ahead and reduce the size of the resistor. The ULN2003 can take it; it's able to sink 500 mA per channel.

Note that the capacitors are rated for twice the voltage of the supply; this is important. And do make sure that you connect the capacitors properly, as they are polarized. If you don't, you'll get a very digital look to the candle, followed by a loud *crack!* of the capacitor popping. I learned this the hard way after a late-night soldering session. What you'll want to do is insert the LEDs into the body so that they just pop up from the well. Then add the capacitor into the control lines so that it can be pushed into its mount hole and the wires routed through the notch, as shown in Figure 3. You can use a bit of duct tape to hold everything in place.
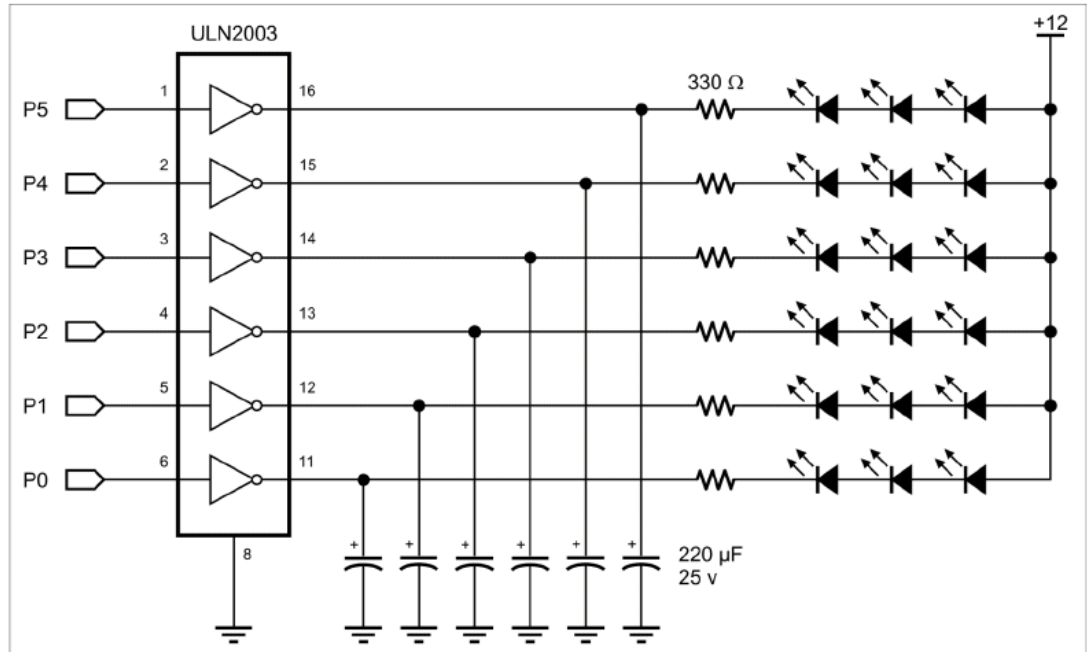

Figure 3. Candle Capacitor Mount.

Candle Code

Okay, the hard part

is done — now for the really easy part: the software. This version is a bit of an update from what I did for Brian a few

potentiometer, the "wind" input is checked. If there is no wind, the flicker rate is scaled to 5 to 15. This produces a nice rate that makes the candle look like it's sitting in a still room. If the wind input is active, then the rate value is scaled to 35 to 60. The longer timing allows the LED capacitor to discharge more completely and creates the illusion that air has been moved around the flame, and that the flame was nearly extinguished.

Okay, how about a bit more realism? Yes, it can be done — if you only need one or two candles (per controller). The idea is to use three independent LED circuits for each wick. This will give movement to the flame, in addition to more modulation of the candle brightness. Use a 510 ohm resistor for the LED (Figure 5), and remember that you have to mount three capacitors on the bottom of the candle and account for four wires.

The code doesn't change dramatically. What we have to modify for this version is the "dark" check; one for each candle.

```
Main:
  RANDOM flicker

Check_Dark1:
  wicks = flicker & %00000111
  IF wicks > %000000 THEN Check_Dark2
    flicker = flicker | %000010

Check_Dark2:
  wicks = flicker & %00111000
  IF wicks > %000000 THEN Get_Rate
    flicker = flicker | %010000

Get_Rate:
  POT RateCtrl, 100, rate

No_Wind:
  IF Windy = Yes THEN Has_Wind
    rate = rate / 25 + 5
    GOTO Flame_On

Has_Wind:
  rate = rate / 10 + 25

Flame_On:
  Candles = flicker & %00111111
  PAUSE rate
  GOTO Main
```

You can see the independent darkness checks at **Check_Dark1** and **Check_Dark2**. If all three LEDs for either candle are out, these routines light one LED to prevent the candle from going completely dark.

One final reminder before I close: Component values will vary, so be sure to run the Pot Scaling program for the editor before you finalize your code. This will let you get the complete range out of the potentiometer.

Now get to it — go build some candles. Remember, candles aren't just for Halloween. They work at Thanksgiving, Hanukkah, Christmas, Kwanzaa, and probably on lots of

other occasions about which I know nothing. Just don't put them on a birthday cake — you'll never get to eat!

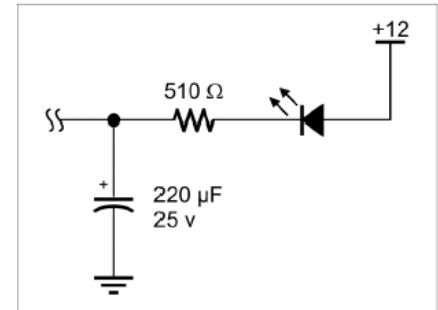Until next time, have a safe and very fun Halloween and, as always, Happy Stamping! **NV**

Figure 5. Using Three LEDs Per Candle.

Electronics Q&A

In this column, I answer questions about all aspects of electronics, including computer hardware, software, circuits, electronic theory, troubleshooting, and anything else of interest to the hobbyist.

Feel free to participate with your questions, comments and suggestions.

You can reach me at:

TJBYERS@aol.com

What's Up:

Magnetic permeability explained and measured. How to test transformer insulation. A crash course in Trigonometry 101, how battery manufacturers test lead-acid batteries, and some very cool Halloween websites.

Permeability Defined

Q. The inductance of an inductor is given by $L = \mu N^2 A / l$

where
μ — Initial permeability of the core
N — Number of turns
A — Cross sectional area of the core
l — Core length

Now, what is the initial permeability? What is the difference between the initial permeability and the effective permeability? Why isn't the average permeability of the hysteresis loop used in the calculation rather than the initial permeability?

M. John

Figure 1 **Field Strength**

Color	Material Type	Permeability
Tan	0	1
Blue	1	20
Red	2	10
Gray	3	35
Yellow	6	8
White	7	9
Black	10	6
Green/White	12	4
Blue/Yellow	17	4
Red/White	25	25
Yellow/White	26	75

Table 1. Iron Powder Core Color Code.

A. In a word: field strength. Actually, that's two words, so let me explain. Permeability is the ability of a material to maintain a magnetic field (air has a permeability of 1) and is defined as the change in magnetic induction (B) for a given change in magnetic field (H). Mathematically, permeability is expressed as $\mu = \Delta B / \Delta H$. As the magnetic field strength increases, so does the permeability — up to a point. When the material is magnetically saturated, permeability peaks, then begins a decline (Figure 1).

Initial permeability describes the permeability of a material at low values of B and may be listed in data sheets as absolute permeability. One definition of initial permeability is 3% of its maximum value. Relative permeability is the ratio between the permeability of the material under test in relationship to the permeability of free space (vacuum) – μrelative = μmaterial / μfree space or $\mu r = \mu 1 / \mu 0$. Effective permeability is often used for cores that have air gaps. This makes the calculations easier because you can ignore the gap by pretending that you are using a material whose permeability is lower than the material itself. Effective permeability is usually relative to initial permeability, and may be listed as average permeability.

Now that you're totally confused, let's assume you're winding a balun coil. It's very unlikely that you're going to come close to saturating the core, so one must assume that the permeability is at its lowest — initial permeability. How much is that? Look to the next question ("Permeability Measured").

Permeability Measured

Q. I am a radio amateur and my aim is to build a 4:1 balun for antenna use, and would like to use some of the toroids I have in my junk box. How do I determine the permeability of toroids with unknown characteristics? My test equipment includes a 200 MHz scope, an MFJ antenna analyzer, and an RF signal generator.

Don Fared K7DRF

A. Ferrites are roughly divided into two groups: iron powder cores and ferrite cores. Iron powder cores are usually color-coded and have permeabilities in the range of 1 to 75 (Table 1). Those with permeabilities up to 800 are usually made from nickel-zinc material and aren't color-coded; over 850, the ferrites are mostly made from manganese zinc.

Using an inductance meter, wind 10 turns of wire equally spaced around the core and measure the inductance (L) in mH. You can then calculate an A_L value for the core using the formula $A_L = 10,000 \times L$, and compare the A_L value to those of known cores of the same physical dimensions listed at **http://amidon-inductive.com/aai_ferriteco res.htm**. Table 2 shows the core types and the permeability of toroids that you're likely to have in your junk box. If you're looking to buy a toroidal core for a project, check out Alltronics (408-778-3868; **www.alltronics.com**).

Don't have an inductance meter? Try the circuit in Figure 2. It attaches to a DMM and reads out inductance in the range of 100 μH to 5 μH. To calibrate, set the DMM to the 200 millivolt range and attach a 1 mH choke across J1 and J2. Adjust the CAL pot so that the DMM reads 1.000.

HiPot Testing

Q. I would like to know if a circuit is readily available for testing the insulation strength of a transformer (HiPot test) at about 1,500 volts according to the equation:

High Pot test voltage = 1000V + 2(mains supply voltage)
= 1000V + 2 x 120VAC = 1,240VAC

B. Bill

Material Type	Initial Perm.	Max Perm.
43	850	3,000
61	125	450
64	250	375
67	40	125
68	20	40
72	2,000	3,500
73	2,500	4,000
J/75	5,000	9,500
77	2,000	6,000

Table 2. Ferrite Core Permeability.

A. The configuration is a simple matter of stacking an AC voltage on top of a DC voltage, as shown in Figure 3. This is a test that is described in detail in MIL-STD-202 (**www.dscc.dla.mil/Programs/Mil Spec/ListDocs.asp?BasicDoc=MI L-STD-202**). And the document stresses the same thing I do: You are working with lethal voltages at times. Always keep that in mind before you throw the power switch!

What you didn't specify is the duration of the applied voltages (typical times range from 10 sec to 1 min) or the environmental conditions. These are variables, too. Higher tempera-

Inductance Tester

Figure 2

NiPot Test Configuration

Figure 3

Load		Battery Aging (10 milliohms initial)				Internal Resistance Voltage Drop
Watts	Amps	100%	80%	67%	50%	
35	2.8	28	35	42	56	mV
50	4.0	40	50	60	80	mV
75	6.0	60	75	90	120	mV

Table 3. Internal Battery Resistance, 27 Series Gel-Cell.

tures, for example, take a toll on insulation. And the longer the voltages are applied at these elevated temperatures, the more likely the insulation is to break down. You need to look at your application. Is this a power transformer operating 24/7, or a furnace control transformer with a lower duty cycle? Are we talking Bangladesh (high humidity)? Iraq (high temperature)? Each test must take into consideration the conditions under which the equipment is to be operating. The HiPot equation – DCV + 2ACV – is just one part of a bigger picture.

Sailor Needs Battery Tester

Q. I read with interest your article, "Pardon my Impedance" about internal battery resistance, and wonder if it may be applicable to my situation. I have a sailboat with a 12-volt electrical system. It uses (2) 8D, (2) 4D, and (3) 27 series gell batteries (I like the concept of no liquid or acid smell) located in different areas of the boat — all tied together in parallel. But the batteries are very heavy and difficult to remove for testing (they only last 6 to 8 years), and the local battery shop doesn't do house

Gel-Cell Internal Resistance Tester

Figure 4

calls. So, if I could check the internal resistance of each battery without removing it from the boat, I think it would be a great help. With batteries this large, what load resistors do you recommend and what voltage should the battery be?

Dennis Baut

A. You've hit the nail on the head. The pure ohmic measurement of battery resistance is one of the oldest and most reliable methods for determining battery usage. As a rule, a 25% increase in initial (newly installed) resistance depicts a 20% reduction in longevity. A 50% increase in internal resistance relates to about 33% of use, and is where most battery dealers draw the line on warranty. My guess is that at double the initial resistance, I'd start looking for a new battery — but a lot depends on plate type and construction (AGM vs. interlaced).

Let's take your 27 Series battery as an example. Most data sheets list its initial resistance at 10 milliohms (0.010 ohms). The best way to test internal resistance is under actual working conditions. I would guess that you never demand more than 35 watts or 75 watts from any one battery on average. That translates into anywhere between 3 and 6 amps.

Now, what is the best way to dissipate that much power using a load resistor? How about a light bulb? Halogen lamps rated at 12 volts are cheap and plentiful — and make excellent load resistors. Table 3 shows the equivalent load current for 35-watt, 50-watt and 75-watt bulbs. Need a bigger load? Add more lamps in parallel and adjust the

current accordingly.

The proper procedure is to disconnect the battery under test from the battery bank. Fully charge the battery. Let the battery settle until it reads about 12.6 volts. Using the circuit shown in Figure 4, measure and record the initial battery voltage. Close S1 for 10 seconds — no more, no less — and record the battery voltage. Subtract that voltage from the initial voltage and compare it to the numbers in Table 3. That's it.

Lightning Strikes: Trig 101

Q. In the September 2004 edition, you discussed lighting strikes and how their location on Earth can be found by triangulation. Could you show me the math on how this works?

Terry

A. Be careful what you ask for — you just might get it. Without good math skills in geometry and trigonometry, you won't get the details. But it's easy to understand the concept if you can visualize triangles. Let me give it a try.

At least 130 magnetic direction finders are positioned across the US to locate lightning strikes. The direction finder finds the location of the strike by sweeping an area and looking for the strongest signal. Anyone who has ever swung a TV antenna or aimed a satellite dish has done radio direction finding.

When the signal is found, a line is drawn on a map from the antenna to the source. But it doesn't stop there. Triangulation uses a second receiver. It, too, calculates the direction from the antenna to the strike. Again, a line is drawn from the antenna to the source, as shown in Figure 5. Where the lines cross is where the strike happened.

While this shows us where the strike happened on a map, it doesn't tell us how far away it was. For this, we need trigonometry – specifically the sine function (the squeamish can skip this part). The law of sines says that if you know two of the angles and one

opposite side, the rest of the triangle can be constructed.

$$\frac{\sin A}{a} = \frac{\sin B}{b} = \frac{\sin C}{c}$$

In our example, angles A and B are 70 degrees and 50 degrees, respectively, giving angle C a value of 60 degrees. The distance between the antennas is 650 ft. If we plug these values into the equation and solve for a

$\sin A/a = \sin C/c$
$c(\sin A) = a(\sin C)$
$a = c(\sin A/\sin C)$
$a = 650(\sin(70)) / \sin(60)$
$a = 705$ feet

That is, the distance to receiver B is 705 ft; solving for b tells us the strike was 575 ft from antenna A.

Whew, am I glad that's over. Back when I learned about sines and cosines, I had to look up the sine values in a book of tables and do the math on a slide rule. Today, you simply enter the angle into a scientific calculator — like the one in Windows. For the math impaired, go to one of the websites below.

www.1728.com/trig4.htm
www.trig.ionichost.com/trig.php

The Hounds of Baskervilles

Q. I own a small house adjacent to a busy city street and cannot keep dogs from defecating on my lawn. Knowing that most commercial products simply don't work, my friend has tried to configure his own high-pitch frequency device to solve the problem. Unfortunately, after many tests, it doesn't work. Do you have a schematic to help me with this?

Alvena Teufel

A. As far as I know — from experience — these sonic devices don't work. And that's probably because of the simplistic way they are designed. Most sonic "keep-away" devices output a continuous or 60 Hz modulated sound of low intensity that most dogs ignore. After all, dogs don't run away from squealing breaks at frequencies that you or I can't hear. I suspect the answer is to make the sound painful enough to ward off the intruder without disturbing the neighborhood. That would be a 92 dB blast of 22 kHz only when your perimeter is breached — not before, and not periodic like a fog horn. Call it condition training.

Here we have a problem, because finding the components to put this together can get expensive. Let's start with the transducer — the speaker. The most efficient would be a marine piezo transducer like the kind they use for underwater sonar. Unfortunately, they can run hundreds of dollars. A good alternative is a high-wattage horn tweeter used in stereo systems, which you can find for under $50.

Figure 5

Triangulation

$$\frac{\sin A}{a} = \frac{\sin B}{b} = \frac{\sin C}{c}$$

Triangulation

Step 2 is an amplifier that can deliver the volume level of a diesel truck running through your front yard — about

1,000 W, give or take according to the efficiency of the transducer. Power amps designed for auto sound systems can provide this level of sound for about $100.

Next, you need a 22 kHz sine-wave oscillator with enough output to fully drive the power amp. The waveform doesn't have to be clean, and can be generated with a first-order filtered 555 astable oscillator circuit.

Finally, you need to trigger the "alarm" only when a dog steps on your lawn. This can get tricky, but if you know the hours that the dogs are likely to trespass, you can simply press a doorbell button to activate it. Otherwise, your choices are PIR, photoelectric, perimeter control, etc.

Have you considered a horn off a diesel truck — the kind that run off compressed air and are REALLY loud? That, too, is condition training — not only for the dog, but also for its owner.

MAILBAG

Dear TJ,

I am writing about GSplit software that you mentioned in the July 2005 issue. There is a similar program which is more versatile. It is HJ-Split, from **www.freebyte.com/ hjsplit** I have used it and it works fine. Versions are available for platforms other than Windows, including Mac, Linux, DOS, and even Amiga. Moreover, files that are split on one platform can be recombined on another. It can also split large hard disk files into parts that will fit on a CD-Rom for hard disk backup.

Bill Stiles
Hillsboro, Mo.

Dear TJ,

As a footnote to your July 2005 answer about back EMF, diodes, and relays, I have been using a standard LED with no current limiting resistor across the relay. You can actually see the inductive kick as the LED winks at you when the relay de-energizes. Yes, I know the forward voltage is higher then a regular diode, but I haven't had any problems with destroyed driver transistors.

Dennis Baut

Response: *This will work with small relays like the kind RadioShack sells, but I wouldn't trust it with industrial relays with large coils. — TJ*

The Latest in Networking and Wireless Technologies

Open Communication

Wireless Mesh Networks — *Networking to Connect Almost Anything to Everything Else*

When you remember that a network is just a system of interconnections that link people or things together, it seems as though almost everything in our lives is networked. Isn't our system of roads and highways just a huge network that links towns and cities together? And what about our electrical power distribution system? It uses high power lines to connect homes, offices, and factories to power plants, and one power grid to another. The telephone system is a humongous network. The cable TV system is a big fat network. Of course, we have come to associate the word network with computers.

Today, what computer is *not* networked? If it is not connected to the Internet via a dial-up, broadband cable, or DSL modem, it is connected to a local area network (LAN), which, in turn, is connected to one or more other networks. And, LANs are not restricted to offices — they are also in many homes.

While most of these networks are wired, copper, or fiber optics, a growing number of networks are wireless. Our cell phone systems are big wireless networks that are connected to the telephone network, as well as to the Internet. Some LANs are also wireless, with a person being able to link his or her laptop up to an access point inside a company or at a hot spot in an airport or hotel. Wireless networks are a growing phenomenon. The latest of these is the wireless mesh, a network with the ability to connect almost anything to everything else.

Mesh Networks

A mesh network is one in which one computer or embedded controller, called a node, is connected to all the others. This is referred to as a full mesh. Figure 1 shows a full mesh of four computers, where all computers have a direct link to all other computers. This is pretty useful because all computers can talk to one another. The concept quickly gets out of hand, however, the more computers you add. The formula below lets you calculate how many individual links (L) you need to connect N computers.

$$L = N(N - 1)/2$$

With 10 computers, the number of links grows to L = 10(10 – 1)/2 = 90/2 = 45. Yikes! It gets messy fast, doesn't it? Can you just imagine all the wires for 100 computers?

For precisely this reason, the full mesh is rarely used. Instead, networks use other methods of interconnecting computers or nodes. These methods are known as topologies. The most popular topologies are the bus and the star (see Figure 2). The bus is a single pair of wires (coax or twisted pair) connected to all computers. It is set up so that only one computer can send to one other computer at a time. The star uses a central computer to coordinate the exchange between nodes. A third, less popular topology is the ring (not shown in Figure 2) that passes messages around a closed loop to the correct destination with only one node transmitting at a time.

Single-node transmission is the key to most topologies. Not all nodes want to talk at the same time, so you can simplify the number of connections and share the common connections. Because of the common connection, you need an access method or software procedure to take control and make decisions when more than one

Figure 1. Full mesh network with wired links.

Computer-Node

Wired link

Figure 2. The most popular network topologies: (A) – Bus, (B) – Star.

Computer Node

Common Bus

(A)

Computer Node

Central control computer

(B)

node tries to transmit simultaneously.

In order to network computers wirelessly, you normally use what is called a point-to-multipoint (PMP) connection, which is simply a wireless star topology. A central computer coordinates the transmission and reception between all other nodes. This works pretty well, but it does not provide a good solution to all applications. An alternative is the wireless mesh.

A wireless mesh is not a full mesh in that, while all nodes can communicate with one another, their paths are not direct. Each node can talk to any other node by using paths through one or more other nodes. This is called a partial mesh and is illustrated in Figure 3. Each node is a computer or an embedded microcontroller — most likely the latter. Each node contains a radio transceiver that can link up with a nearby neighbor. The most common type of node is a sensor or transducer that measures temperature, pressure, light level, a switch closure, or other input. Another type of node contains an actuator, a relay, switch, solenoid, or something that turns a motor or light off or on.

The transmit power is very low so range is not great. In practical sensor meshes, the range between nodes may only be a maximum of 30 feet or so. So, how does any node get a message to any other node? By transmitting the message to a nearby neighbor, which in turn retransmits it to another closer neighbor, and so on until the desired destination is reached. Each node is set up to act as a repeater or router so that the message hops from node to node. In Figure 3, node A can get a message to node H by hopping through nodes B-D-F. An alternate path consists of nodes A-C-E-H. You get the idea.

A mesh network has some pretty cool features and benefits. First, the range of any node is greatly extended because it can send and receive messages over long distances, providing there are nodes between the two to pass the message along. Second, there are alternate paths through the mesh. If one node dies, due to a dead battery or a defect, or if the signal is blocked physically by walls or other objects, not to worry. Likewise, it's no problem when a portable or mobile node moves so that the signal cannot get to one of its neighbors. When that happens, the signal simply finds an alternate path through the remaining nodes.

For example, suppose the battery on node F died and the path between nodes C and E is blocked by a truck. A can still talk to H via A-C-E-H or A-B-D-E-H. We say that the mesh is "self-healing." This feature makes the mesh one very reliable network. To optimize this feature, however, you need to have many nodes to provide those alternate paths. When you need to have a robust wireless network, a mesh is the way to go.

Another desirable characteristic of a mesh network is that it is self-forming. In other words, the nodes are designed to automatically discover one another and connect. This happens as two nodes fall within range of their transmitters and receivers. They exchange hellos and a link

Figure 3. A partial mesh network with wireless links. It extends the range and makes communications more reliable.

is formed. This type of network is typically called an ad-hoc network. Mobile or portable nodes may come and go, so the network is constantly reconfiguring itself to the remaining nodes. Nodes can be quickly added or removed without much consequence and certainly with no need for human intervention.

An important consideration is that most mesh nodes are portable or mobile, meaning they are battery-powered. Because they use very low power, batteries can last for ages. Most nodes go into a sleep mode until they have data to send or until they are called upon to relay another message. They are probably on for less than 1% of the time. This very low duty cycle and the low current drain during the sleep mode (only about 1 μA) means that a battery can have a life up to a year or more. For industrial applications, battery replacement and maintenance costs become very low.

Mesh Applications

As with many electronics, first the method is discovered and then we sit around and figure out what to do with

Figure 4. Ember's ZigBee developer's kit.

it. That is happening right now with mesh applications, but many great uses have already been discovered and implemented.

The military was probably the first to use a mesh concept. Each soldier is equipped with a mesh-enabled radio that can link up with the radios of other nearby soldiers. All can communicate with one another and messages are easily relayed to more remote nodes through the mesh. The mesh idea can be expanded to include weapon systems of all sorts, giving commanders a way to monitor and control the battlefield despite constantly changing conditions, such as units moving great distances away.

One growth application is building automation. In office buildings, hotels, hospitals, and universities, mesh can ease the process or monitoring temperature, lights, HVAC, and other conditions, and taking control steps to improve efficiency. Hundreds of battery-powered nodes can be interconnected wirelessly to mesh an entire building very quickly and easily. Just think of the cost and time savings of having to wire all those sensors and controllers. Mesh makes it possible to achieve the same results wirelessly.

Home monitoring and control is another growing use of mesh. Again, you can use "stick-on" nodes containing all sorts of sensors and controllers for light switches and appliances. Again, the big benefit is that there is no wiring or configuring. Just put the nodes in place and let them automatically link up and get to work.

Industrial monitoring and control is another growing use of wireless mesh. Industrial wiring is horrendously expensive, from $200 to $1,000 per foot if run in conduit by a licensed union electrician. One sensor wire can cost you up to $100,000. Ridiculous. But with wireless, it is far less expensive to put sensors on everything to monitor process conditions and take control actions. In manufacturing, mesh is being used to monitor the well-being of the various machines in order to predict their failure and to determine when to perform preventative maintenance.

Wireless automatic meter reading is taking off as a great application of mesh. Reading electric, gas, or water meters is expensive and time consuming. But with mesh, it is fast, easy, and inexpensive. And the mesh from house to house makes it possible to collect all the data from remote points in the neighborhood.

A variety of other uses are being tested and developed. Environmental monitoring of forests, beaches, and other sensitive areas is easy with wireless sensor networks. Medical monitoring of patients in a hospital is another use. Security systems can use a greater quantity of wireless mesh nodes because they are cheap and battery powered. An upcoming application is using mesh to eliminate wiring in home entertainment systems. Mesh networks using UWB (ultra wideband) can transmit at speeds over 100 Mbps, making it easy to connect video devices like HDTV sets, DVD players, cable boxes, and the like without wires. Look for these next year.

How to Make a Wireless Mesh

The RF transceiver part of the mesh nodes is simple. There are literally hundreds of different chips and modems available. What makes a mesh a mesh is the software embedded in the interconnected controller. Almost any air interface can be used. Simple designs can use the ISM band ASK or FSK chips operating in the 315, 433, and 915 MHz bands. Bluetooth is not as useable, but could be utilized in simple cases. Bluetooth inherently contains an ad-hoc network of the star or point-to-multipoint nature. It can automatically link up with seven other nodes forming a piconet. And the controlling node can link up with other controller nodes forming scatternets. While this may be suitable for some simple applications, it is not a mesh.

A better choice is ZigBee, a standard designed especially for mesh applications. ZigBee is the marketing name for an expanded version of the IEEE's 802.15.4 standard. This standard defines the radio interface and the data link or MAC (media access control) method of handling

Figure 5. The MOTOMESH network for public safety and service. The access points are designated EWR and meshed so that any remote node can talk to any other node.

data in and out.

The ZigBee radios can work in the 868 MHz (Europe) or the 915 MHz and 2.4 GHz bands (US). The mode of operation is offset-QPSK using direct sequence spread spectrum. Data rates are 20 kbps for the 868 MHz version, and 40 kbps for the 915 MHz version and 250 kbps version. Most applications have opted for the 2.4 GHz 250 kbps version. Yes, the data rate is low, but most applications only generate very low speed digital data, like that from a temperature or pressure sensor or a switch closure, so high-speed connections aren't needed.

The ZigBee Alliance, an organization that promotes ZigBee as well as develops additional software to make mesh work, also conducts testing and certification of ZigBee products to make sure that they all interoperate. For details, go to **www.zigbee.org**

There are dozens of companies gearing up to offer ZigBee chips, software, and systems that make it easy to create a mesh network. One offering the most complete solution is Ember. They provide chips, boards, software, and a development system pictured in Figure 4. Note the 12 node boards you can use to set up a representative mesh to try out all the various options. Check out **www.ember.com**

Another popular wireless technology that can be used to form a mesh is the widely used wireless LAN standard 802.11 also known as Wi-Fi. Most laptops use this to link up at available hot spots for Internet access and email. Companies and other large organizations with many networked PCs use wireless access points to extend their existing wired LANs.

There are three primary problems with 802.11. First, it was not designed for mesh use. It is inherently a point-to-point or point-to-multipoint technology. Yet, it can be made to do mesh networking with additional software. Second, it is more power hungry than ZigBee. Battery operation is not practical, except for short intervals of a few hours or less. If you are near power, then this is not a disadvantage. Third, the range is much greater than the

typical mesh node, which is rarely more than 30 meters. With Wi-Fi, you get a range up to 100 meters. The good news is that you have access to fast data rates, 11 Mbps in the basic 802.11b system and up to 54 Mbps in the 802.11a/g standards.

One great example of using Wi-Fi in a mesh is creating a broadband access service for those people who do not have a high speed Internet connection via cable TV or a DSL line. This covers many folks who live outside a major metropolitan area or in a rural district.

By giving each subscriber a Wi-Fi-enabled transceiver with special mesh software, each node can then talk to other nearby nodes and relay data for others as well as receive and transmit data. A nearby access point collects all the inputs and provides a connection to an Internet service provider. This is already being done in many areas around the country.

Another mesh using technology like 802.11 is Motorola's MOTOMESH. This sophisticated mesh uses access points called EWR (Enhanced Wireless Routers) that incorporate both 802.11b Wi-Fi radios operating in the 2.4 GHz band and two similar radios operating in the 4.9 GHz band. This 4.9 GHz band

is allocated for public safety systems used by police, fire, EMS, public works, and other city and county services.

The EWRs are all meshed together as shown in Figure 5. Portable laptops can link into an access point and talk to any other laptops. The MOTOMESH uses Motorola's Mesh Enabled Architecture, a system originally developed for the military. With data rates to 6 Mb/s, it can enable high-speed data transmissions as well as video and graphics. And it can do this while traveling at speeds to 200 mph. And a neat feature of the MOTOMESH is its built-in location technology that lets any node determine the location (latitude/longitude) of itself or any other connected node. For details, see **www.meshnetworks.com** or **www.motorola.com**

The MOTOMESH can also be used for telemetry applications. It was recently used by Chevy's Corvette group at the American Lemans Series races where Corvettes took first and second place. Part of the victory was due to the team being able to predict the need for pit stops using the telemetry data from the racecars. This helped eliminate some pit stops altogether and time the others.

If you are not using mesh now, you soon will be in some form. **NV**

SMB INTERFACE FOR BOARD-TO-BOARD APPLICATIONS

Bomar Interconnect Products, Inc., a premier manufacturer and designer of quality connectors and accessories for RF, video, and broadcast transmission, announces the addition of an SMB interconnect to their proprietary Eliminator™ Series. The coaxial Series, which already consists of a family of F connectors, is expressly engineered to eliminate the need for extra cable assemblies and PCB connectors traditionally required to mate box-type or modular integrated components to boards. Manufacturers need only connect an Eliminator directly to an electronic device and drop it onto the PCB. In addition to significantly reducing installation time, the subminiature SMB Eliminator promotes board space conservation along with circuit miniaturization. The interconnect is available in 50- and 75-ohm impedance, and provides broadband performance with extremely low reflection

DC through 4 GHz.

Bomar's newest right-angle interface is ideally suited for employment in multi-board applications in which digital to RF conversion is required. As such, it may be confidently specified for use on PCBs in a broad range of military, industrial, consumer, and medical equipment.

Quality engineered of precision machined brass, the new SMB Eliminator features TFE insulators and gold-plated contacts to ensure continuous reliability. Additional features include an insulation resistance of 1,000 mOhms min., a dielectric withstand voltage of 1,000 VRMS, and a temperature range of –65 degrees C to +165 degrees C. The Series is fully customizable to accommodate various PCB thickness, plating options, and board-to-board spacing. A die-cast version is also offered for high-quantity, low-cost applications.

Pricing for Bomar's SMB Eliminator starts at $1.38 each in quantities of 1,000. Delivery is from stock to 12 weeks.

For more information, contact:

BOMAR INTERCONNECT PRODUCTS, INC.
1850 Route 46 E.
Ledgewood, NJ 07852
973-347-4040 Fax: **973-347-2111**
Email: **mail@bomarinterconnect.com**
Web: **www.bomarinterconnect.com**

BABY ORANGUTAN ROBOT CONTROLLER

Pololu announces the release of the Baby Orangutan robot controller, the latest addition to Pololu's line of Orangutan robot controllers. The Baby Orangutan features an Atmel mega168 microcontroller with 16K of Flash program memory and 1K data memory, two H-bridges for control of two independent DC motors or one stepper motor, two LEDs, and a potentiometer connected to an analog input. The unit is compatible with the larger Orangutan robot controller, though some features — such as a display — must be connected externally.

The compact module has dimensions of 1.2" x 0.7", and it can be configured to fit in a breadboard or a 24-pin dual in-line package (DIP) socket. For applications with low I/O usage, the Baby Orangutan board can also be configured with pins on just one side of the module for use as a single in-line package (SIP). The diminutive size of the Baby Orangutan makes it well suited for primary control of miniature robots or for auxiliary control on larger robots.

The Baby Orangutan is available with two

programming connector options: a standard, 0.1"-grid header for the greatest convenience, and a low-profile, flat flexible cable (FFC) connector that reduces the module thickness to 0.110".

The unit price ranges from $24.95 to $29.95, depending on configuration, with free US shipping available on most units.

For more information, contact:

POLOLU CORPORATION
6000 S. Eastern Ave. Suite 12-D
Las Vegas, NV 89119
877-7-POLOLU or **702-262-6648**
Fax: **702-262-6894**
Email: **www.@pololu.com**
Web: **www.pololu.com**

EMBER WIRELESS MESHING SOLUTION GOES MODULAR

Recently, at the ZigBee Alliance Open House in Oslo, UK-based radio specialist Telegesis announced the arrival of ETRX1 — a meshing RF module that enables easy integration into industrial and other applications. Designed as a complete hardware solution on one tiny board, the ETRX1 will allow companies to quickly tap into the lucrative ZigBee market for wireless control and monitoring products without requiring in-house RF engineering experience.

Incorporating the ZigBee ready Ember meshing stack — EmberNet — the ETRX1 module is built around Ember's EM2420 radio and an Atmel Atmega 128L microprocessor. An over-the-air upgrade to the new, fully compliant EmberZnet stack will be available within a short time.

In addition to the compact hardware design, Telegesis has introduced a comprehensive AT style command layer to facilitate rapid adoption of the technology. This familiar format will enable engineers to communicate easily and simply with the meshing software and move quickly to full integration.

The module measures only 37.75 mm x 20.45 mm and is surface mountable. It will carry a GigaAnt RUFA™ antenna and will be fully CE/FCC tested, eliminating the need for end-product certification. Alternatively, customers can specify a Hirose H.FL connector for external antenna applications. The ETRX1 operates in the 2.4 GHz ISM bandwidth and the technology is IEEE 802.15.4 compliant.

To compress development times, OEMs can assess the ETRX1 using the supporting development board which is available in a low-cost kit complete with power pack, RS232 cable, three modules, and downloadable Telegesis Terminal software.

For more information, contact:

LEMOS INTERNATIONAL CO., INC.
(North American Distributor)
Fairfield, CT USA
Email: **Sales@LEMOSINT.COM**
Web: **WWW.LEMOSINT.COM**

SMT MICRO LEAD FRAME ADAPTORS

Bellin Dynamic Systems announces the latest Snap-Apart™ adapter boards: the P532 and the P533. The P532 and P533 are designed to adapt SMT Micro Lead Frame devices to the .100" standard.

SMT chips can be difficult and frustrating for electrical engineers to prototype. The P532 and P533 can quickly take QFP, MLF, QFN, or any comparable package and place them in a familiar DIP .600" footprint. By utilizing the correct IC, engineers can keep board revisions to a minimum.

Both adapter boards hold eight Snap-Apart™ elements and feature eight Universal adapters. Each board has four adapters that can adapt devices with up to 80 pins and four that can work with devices up to 32 pins. They work with quad devices that have .4 mm pitch and/or .5 mm pitch. The P532 is designed for quad pack-

ages with an even number of pins per side. The P533 works with quad packages with an odd number of pins per side. The elements easily snap apart from each other to provide individual adapters. Each P532 and P533 board comes with six 40 pin header strips for quick and economical prototyping.

Bellin Dynamic Systems is a provider of rapid prototyping tools for engineering development. Visit their website for data sheets and more information on this and other SMT solutions.

The P532 is currently available for $59.95 each. The P533 is currently available for $59.95 each.

For more information, contact:

BELLIN DYNAMIC SYSTEMS, INC.
714-630-8024 Fax: **714-630-8025**
Web: **www.beldynsys.com**

LOW-COST IR TRANSMITTER CHIP

A new trainable, low-cost IR transmitter chip offered by Tauntek can make controlling other equipment from your next project a breeze. It works with most equipment and remotes and can learn up to 16 commands — each of which can be a different format/protocol. It has a keyboard or MCU interface and it is based on the PIC16LF87 Micro with EEPROM. The chip alone, PC boards, or complete kits are available. Please send inquiries to the email address below.

For more information, contact:

TAUNTEK
Email: **rlgrieb@verizon.net**
Web: **www.tauntek.com**

by Mike Keesling

Understanding, Designing, and Constructing Robots and Robotic Systems

Personal Robotics

The Ball Bot

Ugh! It's column time, and my editor has been after me to do some coverage on the Ball Bot, something I had started a while ago, but have not yet completed.

To recap, a simple description of the Ball Bot would be "a two-axis inverted pendulum." The operative word here is simple. In its implementation, it is three belt-reduced, encoded motors driving some Omniwheels sitting on top of a basketball — with a bunch of inertial sensors and a DSP thrown in for good measure.

Originally, I had several different thoughts on how to accomplish this build, but I have managed to narrow the concept down quite a bit. Originally, the scope of my build was massive. Rather than rant about in my usual fashion, I will approach my thought processes methodically.

The Ball

Originally, I fully intended to manufacture my own sphere. I wanted to control the diameter, resiliency, surface, appearance, and all of that good stuff. I have decided to scale back any grandiose notions I may have had.

(Insert pitiful excuses here.) My garage is sweltering hot and my house is nearly unbearable this summer. The last thing I want to do is wrap myself in a cotton jumpsuit and breath mask and start pouring, cutting, grinding, shaping, molding, and forming my ideal of a perfect sphere when all I really need to do is hop on over to the sporting goods store and buy one. In all fairness, making a mold and then a casting is a time consuming, expensive business. I am not precluding it as an eventual possibility, but I'll be more receptive to the idea this winter.

Now, I must say that I did get stares from several sporting goods salespeople. There I am in the aisle, sitting cross-legged, scrutinizing every orb, seeking out every defect and every inaccuracy. I was poking, prodding, pushing, squeezing, and bouncing.

After my intense examination, I decided on a basketball. Yup, a basketball. A good ol' men's regulation-size 7, 29.5" circumference, 9.39" diameter basketball. It's not in the least bit glorious, but it is frugal, and I will convince myself one day that there is just as much panache in simplicity ...

Omniwheels

I am still sticking with Omniwheels. I can see no other way to generate the motions that I need on the surface of the sphere. Kornylak (**www.kornylak.com**) has a great selection of Omniwheels. I went with a small diameter, because I plan on making the robot on the shortish side, and will therefore need some good acceleration to get "back under" the main mass of the bot. I can get the speed I need in other ways.

The one issue with all the Omniwheels I have seen is the challenge of linking them to the shaft they ride on. I still haven't conquered this one. The Omniwheels have a sort of fluted end to which I am hoping to engage the shaft, via a machined hub, and possibly a keyed shaft. I will cross that bridge when I get there. Kornylak does offer some wheels with a chain drive hub that locks to them on their shaft, and even allows you to clock their rotation via multiple keys, but I wanted to suffer and use belts.

Drive Train

I went to my stash of motors and

found some nice encoded coreless motors. They are 27-watt motors, rated at 14.3 oz-in stall torque, 5 oz-in continuous torque, 2 amps continuous current, and 6,900 rpm at 24 volts — a beefy motor to be sure. With my planned 4:1 primary belt reduction and the reduction from the small 2.58" Omniwheel driving the basketball, I anticipate around 300 rpm at the wheel after you subtract the speed loss from torque. I believe there will be some speed limits that I have to accommodate for, based on the fact that most of the time, a wheel will not actually be contributing all of its motion to moving in a vector perpendicular to its axis of rotation.

The one bugaboo here — besides coupling the Omniwheel to the shaft — is the diameter of the shaft itself. It is 8 mm, and that makes bearings a bit pricier and a bit tougher to find. The outside diameter of the bearing is 22 mm, which also makes construction tougher. I was looking at appreciable sums just for bearings, but an eBay search for "22 mm flanged bearing" resulted in cheap satisfaction.
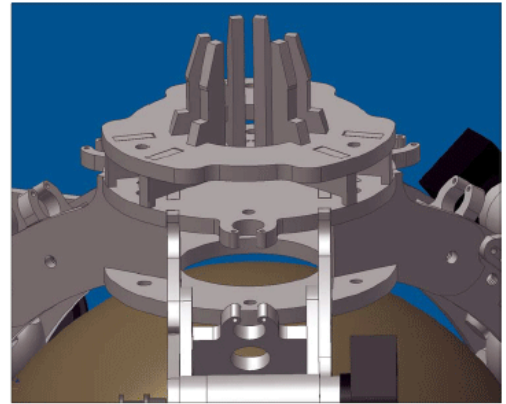
Construction

I opted for water jet cut aluminum with tabbed and grooved construction. It's very similar to how I build laser-cut acrylic things, but with bolts and tapped holes holding it all together. With nothing welded, if I make changes, there are no tears to shed. I may well have to accommodate a bigger hub, for instance, which will necessitate some new parts. Experience has taught me about kerfs on water jet material, so everything fits together just right. No filing, just some tapping.

My choice of water jet cutting was simple, as we have one at my new place of employment. 'Nuff said. I certainly wouldn't tackle this sort of thing unless I knew I could crank out parts pretty readily. Alternately, I have to bet you could throw something together with plywood, marine "T-nuts," all-thread, and polyurethane glue.

"Plywood?" you say. Yes, plywood. The stuff they build houses and boats out of; the stuff they used to make airplanes out of. Wood is a viable alternative to metal, is easy to work, and is strong and inexpensive. I spent $56 on aluminum. I could get a lot of plywood for that much. "Allthread?" you say. Yes, it is really, really poor quality. I have no excuses, other than it is really convenient.

Building entirely out of flats with only perpendicular cuts is a challenge at times, but a fun exercise nonetheless. Clamping the vertical post was especially interesting. I came up with an arrangement of jaws that are held captive, and that are caused to pivot and clamp, by a single ring that is tightened down. This action will hopefully provide enough clamping force to grab the post securely, and if it proves viable, I will replicate it on the opposite end of the post to allow for easy removal of the top section. If this doesn't work, I have the ability to use a hose clamp
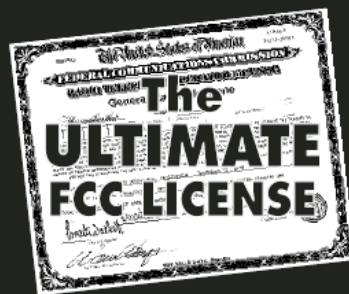

Detail of the clamp assembly.

to augment the clamps.

Sensing and Control

Sensing and control is one of those, "I'll cross that bridge when I come to it" sort of things. The problem is, I can see the bridge, it's coming fast, and I can't avoid it. There are some things I know I must embody, some

Personal Robotics

A possible use of a ball caster to help support the mass.

things I want to embody, and general good practices are involved in doing so.

Spark Fun Electronics (**www.sparkfun.com**) has some nice little inertial sensing boards with a gyro and accelerometer built in. I have a couple, and so far am quite pleased with them. The key to any balancing robot is knowing your gravitational reference. Accelerometers alone would do fine, but they are susceptible to accelerations other than gravity. To tame them, you tend to lower their frequency response. The problem is that then you don't have enough response to sudden forces coming in. This is why you add a gyro.

One may ask, Why not skip the accelerometers and integrate the gyro? Well, it comes down to noise floor. If you integrate rotational velocity and noise into position, you eventually accumulate errors, and you still do not have a gravitational reference. The converse, differentiating your accelerometer, still leaves you with unwanted signals mixed in. It is just easier to divide the functionality into two devices.

The basic concept of the algorithm is to take your accelerometers, which provide tilt, and multiply them by some gain. Then you take your rotational velocity and multiply that by another gain coefficient, and then you sum these into a control signal. This will probably be good enough, but it may be desirable to use a transfer function between the signal and the output. I am sure you could take this to extremes, and calculate the center of your inertia, rather than your tilt

angle, but I will cross the math bridge after I fall into the creek several times.

Once you have a gravitational reference, you still need to move the wheels. Some friends of mine are working on a two-wheel balancer, and have had great success by running the wheels open loop. I am a little afraid of this, however. My wheels must couple to a rubber surface through friction, and must execute coordinated movement in order to produce controlled results, or unwanted/unknown vectors of motion could arise. Additionally, too great of an acceleration could arise in wheel slip.

To that end, I plan on using a beta release of the new MoCon PID motion control implementation from the IsoPod line of microcontrollers. To be honest, I have a bit of a stake in this because I have been acting as a consultant for them. I truly believe that they have a superior product, and that the robotics community needs their stuff. To that end, my efforts have helped them develop a superior product that will soon be at our disposal.

For those of you unfamiliar with PID, it is a way of controlling things to keep them where you want them. PID stands for Proportional, Integral, and Derivative. It is a means of looking at different types of errors, multiplying those errors by gain factors, and setting a new output to your device based on those calculations. While it is most useful to us roboteers for controlling motors, it is also widely used in controlling various industrial processes, from ovens to paper making.

The MoCon word set, as presently implemented by New Micros, Incorporated, sets up an interrupt driven loop with motion control specific words that you stuff into it. It's all very modular, like a PID construction set. Words that calculate PID output based on desired position vs. actual position; words that generate trajectories so you can control position, acceleration/deceleration, and velocity; and words that allow you to easily read quadrature encoders — up to six of them, in fact — on a single IsoPod.

While this all sounds pretty nifty,

it really is the tip of the iceberg. MoCon is more than a PID controller. For starters, everything is modular. Feedback isn't only from a quadrature encoder, it can also come from an analog channel, CanBUS, SPI, or any other data source. You can set up a quadrature encoder, potentiometer, light sensitive resistor, humidity sensor, or whatever as your feedback.

The output is the same. Typically, one would use one or two of the 12 PWM generators or one of the half dozen timers as an output, but you can just as easily use SPI, for instance.

The trajectory generator, another set of words, is just as flexible. Even if you aren't using the PID portion of the code, and even if you are running typical R/C servos in an open loop mode, you can still have smooth acceleration at your command by just using the trajectory generator.

The PID and trajectory calculations, computed as 36-bit intermediate integers, and executed with lightning speed, happen in about 25 uS. This means that you could easily run six quadrature-encoded motors on a 2 KHz servo cycle with plenty of time to spare. The PID calculations, include the usual Kp, Ki, Kd coefficients, and also have a ds term for sub sampling rate, as well as a scale factor to allow for a massive variety in how you implement your gains and feedback.

All of this gets stitched together in a separate loop that runs at the frequency you prescribe. Remember, this is all built into the language, so you can mix and match as you see fit. You could conceivably use the loop for other things as well, such as reading some sensors, or not use that loop at all, and use the MoCon words in other creative ways.

What all this means to me is that I can monitor my sensors, do the math to get a vector and velocity, do the trig to divide the task among three motors, and run those motors precisely. It also allows me to easily convert everything to real world units by separating the output of the sensors and the inputs to the motors.

If you think about the mechanical layout, you will see that speed control is critical. I need to limit the top speed of the motors, and I need to control them precisely. For example, a motion that is parallel to one of the motors will be at an angle to the other two. They each contribute a vector on a spherical surface. The speed they are able to produce needs to be precise, otherwise you will weave around — or worse.
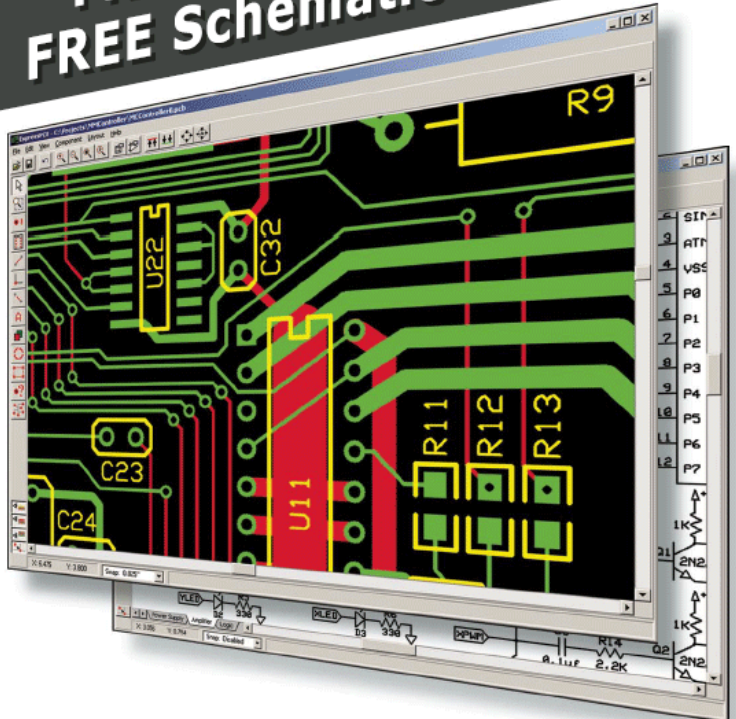
Form Factor

I am really up in the air on this one. Half of me wants to emulate Nomad, from the original Star Trek TV show, and the other wants to do a table. This is part of the reasoning for making a clamp-on arrangement. While my sensing and lighter components will be on the base, my chassis and considerable

mass in batteries will be up high, buried in whatever form factor I decide upon.

Other Considerations and Tidbits

There are a lot of little things that are going to make this project, and many factors to consider.

It will initially be radio controlled. With three quadrature encoders for feedback, I still have 4+ timer pins to interface to a four-channel radio. Presently, I am researching a four-channel terrestrial transmitter and receiver. Remember, it's not nice to fly aircraft frequencies on a ground- or water-based craft.

I have something my friend Phil Davis suggested, called a ball caster, to rest between the ball and the chassis. If I find that the ball gets "sucked up" into my wheels, I can add the ball caster to prevent this — at the expense of drag and increased low-slung mass.

One thing to remember when building an inverted pendulum is that mass is your friend if properly placed. If the proportions of mass distribution aren't correct, you will topple the upper portion of your robot, rather than driving under the side that is falling and catching it. It is like balancing a broomstick vs. a pencil on your finger. With your mass up high, you can position your chassis under it.

To affect motion, you simply tilt yourself and then try to keep up with it. Think of it as a controlled fall.

Another interesting note to consider is motor drive voltage. Torque relates to the magnet flux within the motor. This is the product of voltage and current. As long as you don't exceed the dielectric constant of a motor, or the current capacity of its commutator or wires, you can over drive a motor in short bursts.

The real issue is total watts going into that motor over time. That will build up heat, and that will destroy a motor. If you expect a motor to be driven at a low duty cycle during most of its use, with occasional bursts, you can use a higher bus voltage. This works out well for me, because I can add more batteries, run a higher bus voltage, maintain safe overall duty cycle, and have more mass to put up top.

I plan on running a whole lot of battery. Originally, I considered a string of 30, 3,000 mAh NiMh cells in series, but I may consider five smaller gell cells in series. I am going to avoid LiPoly batteries this time, due to the cost.

Remember, this is all very flexible. I am never afraid to ditch something if it just isn't working out. Don't be surprised if things look different next time I report on this beast. The winter months are approaching, and who knows, maybe next time I cover this project I will have found the fortitude to build my own sphere. Until then, keep rolling. **NV**

"On eBay, I found the exact oscilloscope we needed for our mass spectrometry project. I bought a used one on eBay from J.C. for just $1,500 — instead of the usual $10,000. That means $8,500 more for research."

Noah Peter Christian, Ph.D., Indiana University – eBay ID: noahskoah (41) ⭐

"I've sold electronics since 1992 and started selling on eBay in 2000. With eBay, I can reach specialized equipment buyers around the world — *and* boost sales by $100K each year."

J.C. Christen, Santa Clara, CA – eBay ID: granite197 (534) ⭐

Every day on eBay, great deals are made, new customers are found and business dreams fulfilled. Make eBay a part of your success story.

www.ebaybusiness.com

FREE Get customized selling help. **A $500 value!** **Call 1-866-610-3229.** M – F from 9 a.m.– 9 p.m. EST. Be sure to mention code IET P111.

ebaY ®
business

Keep your friends informed & entertained.

They'll thank you all year around!

Seasons Geekings

(And they won't need to borrow **your** issues anymore)

100% pure electronics

Gift Subscription

ROBOTICS MICROCONTROLLERS COMPUTER CONTROL SPACE
Everything For Electronics
Nuts & Volts
Excellent Circuits With Excel
USB 16-bit AD/DA Converter
SPOKE SIGNAL
LEDs On

ROBOTICS MICROCONTROLLERS COMPUTER CONTROL SPACE
Everything For Electronics
Nuts & Volts
September 2005
CAPACITORS
The Family Tree
Build a Pool Timer
Generating Analog Waves From Digital Signals
The Space Elevator

Simple Circuit Breaker

Don't Waste Time Replacing Regular Fuses — Use This "Electronic" Fuse Instead

The Fuzzball Rating System

To find out the level of difficulty for each of these projects, turn to Fuzzball for the answers.

The scale is from 1-4, with four Fuzzballs being the more difficult or advanced projects. Just look for the Fuzzballs in the opening header.

You'll also find information included in each article on any special tools or skills you'll need to complete the project.

Let the soldering begin!

Even though there are fuses in the car, resetting a circuit breaker is a lot easier than buying and replacing a "one-time-use" fuse. This circuit is simple, uses a handful of readily available parts, is completely electronic, and has a very fast cut off time. A picture of the completed circuit breaker and remotely located FET/Rsense is shown in Figures 1, 2, and 3. Also shown is the completed project PC board. This application is used in a high-amp, low voltage DC supply.

The heart of the "electronic fuse" is the Norton Quad op-amp, LM 2900/3900. This op-amp is unique in that it is current operated and uses a current mirror to furnish the non-inverting input. In addition, it is designed for a single supply power voltage from 4 to 36 volts with only small changes in performance characteristics over the entire voltage range. The op-amp is ideal for this project since wide variations of primary power may be encountered, especially with a transformer/diode bridge/capacitive filter power supply.

One section of the op-amp is used as a comparator to determine when the set current limit is reached. The second section is used as a voltage doubler for the gate control of the N Channel HEXFET (IRF 1404) used as the power switch.

Referring to Figure 4 — the schematic of the project — section B of the LM3900 is used as a comparator with hysteresis to ensure the output remains in a high state until the short circuit is completely removed and the comparator is reset. The current into the non-inverting input (pin 1) is compared to that flowing in the inverting input (pin 6). R15 (25K pot) is adjusted so that, in normal operation, more current is flowing into pin 6 than pin 1.

Once the current limit has been reached, the voltage drop across Rsense is great enough to decrease the current flowing into pin 6 to the point where more current flows into pin 1 and causes the output at pin 5 to go high. Removing the short circuit will not allow the comparator to switch back because of the hysteresis introduced by R6. It is necessary to ground out some of the current going to pin 1 by pushing the reset push button switch. R3, the 3.9M resistor, is not enough to reset the circuit, but the parallel .22 mF cap across it, C6, will draw enough momentary current to reset the comparator and put power into the load again. If the short circuit

Figure 1. The author's low voltage, high amperage power supply. The circuit breaker PC Board is mounted sideways on the lower left of the picture, while Rsense and the FET is located on the upper right of the figure, just to the left of the large filter capacitor. It is mounted on the back panel of the chassis.
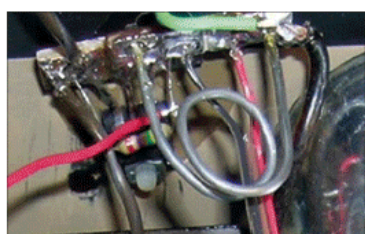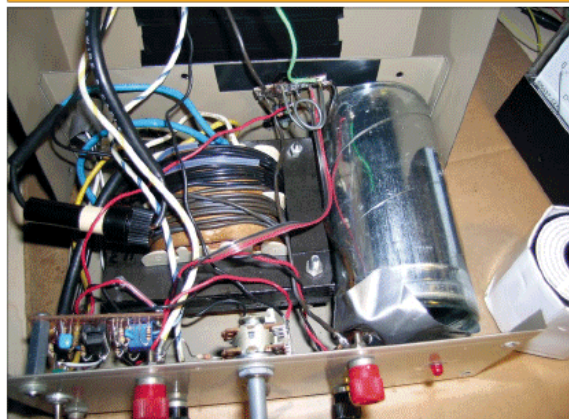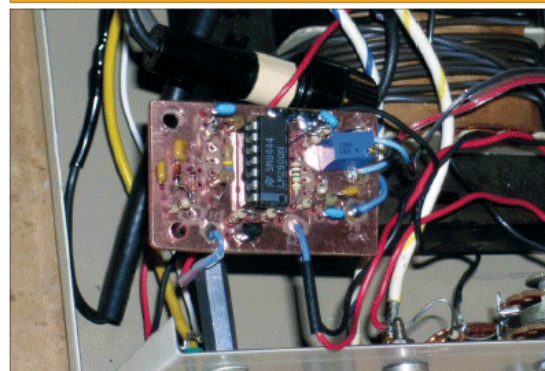

Figure 2. Close-up of Rsense (single coil of steel wire), and the IRF1404 FET. Also shown is R10 mounted on the FET.

Figure 3. Close-up of Circuit Breaker PC Board.

has not been removed from the load side, the comparator will immediately switch on again and stop any current from reaching the load. It is not possible to hold the reset switch closed and continue supplying current to the short-circuited load. R3, the 3.9M resistor, will drain the charge across C6, preparing it for the next reset operation.

Section C of the LM 3900 is used as a voltage doubler, doubling the supply voltage of the LM 3900. That voltage is necessary to turn on the N channel FET completely. A P channel FET could also have been used, but N channel HEXFETs are less expensive, have less on resistance (Rds ON), and are more readily available. The FET used in this circuit (IRF 1404) is rated for up to 40 volts across the source-drain junction, 162 amps, with an on resistance of .004 ohm.
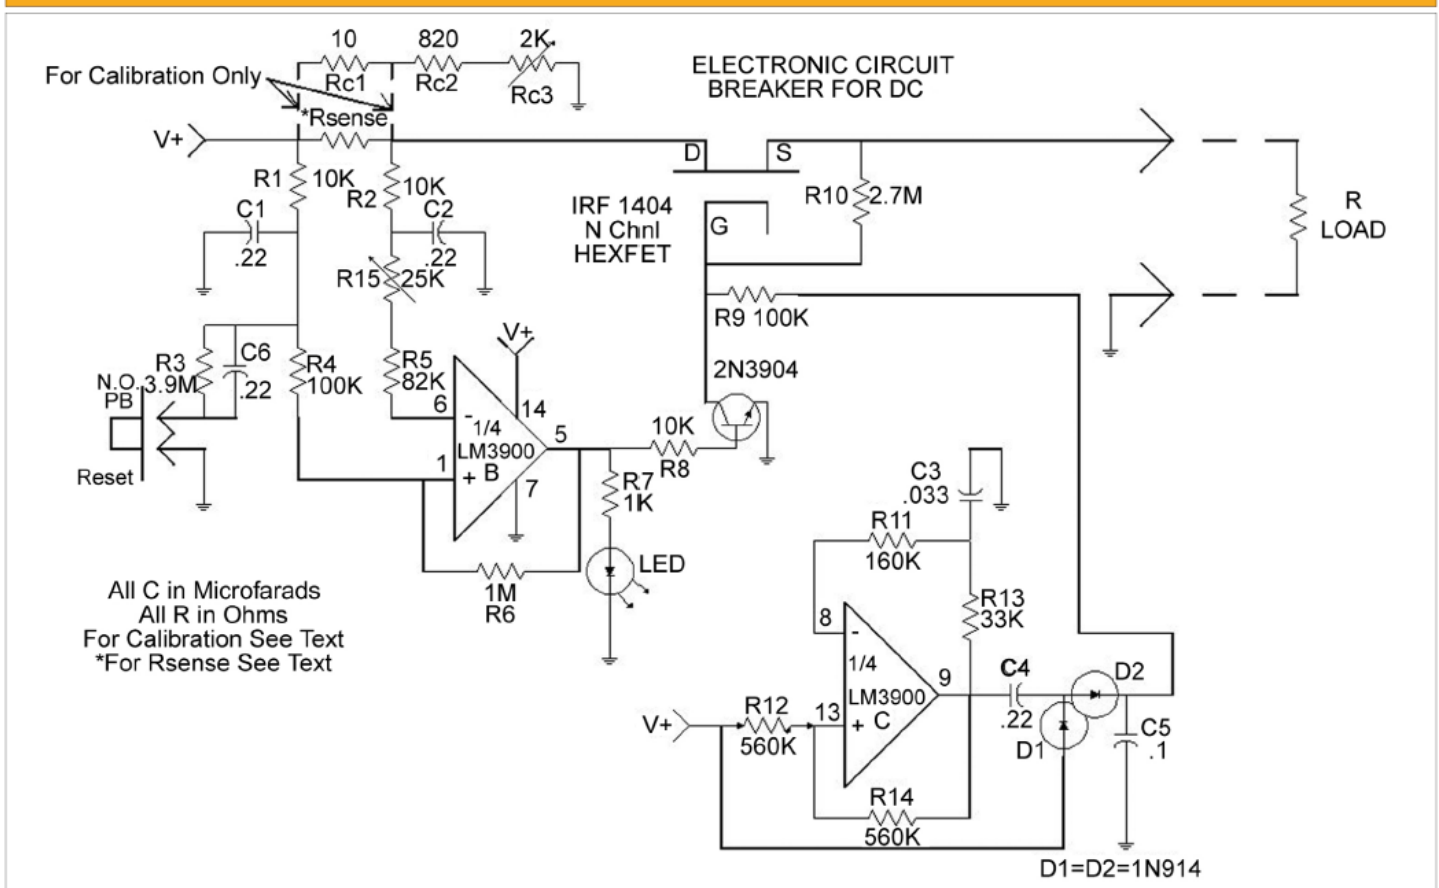
The square wave output from pin 9 of the LM3900 is routed into the voltage-doubling network of D1, D2, and C5, which then connects to the FET gate via a 100K resistor (R9). Also connected to the FET gate is the collector of an NPN transistor (2N3904 or equivalent). Once the comparator has been tripped, it puts out a high on the base of the transistor via R8, turning on the transistor and shorting the drive voltage for the FET to ground, turning it off. A circuit "trip" indicator is provided by an LED also connected to the comparator output (pin 5) via R7. Once the short circuit is removed and the reset button is pushed, the comparator output on pin 5 goes to ground again, shutting off the tran-

sistor allowing the doubled supply voltage back to the FET gate, turning it on again and restoring power to the load.

Rsense is one of the most important parts of this project. It must have the smallest possible resistance, be able to handle large currents, and not be too physically large to handle. I settled on a value of .01 ohm, which will only drop .3 volts at 30 amps. The amperage range of this project may be set between 1 and 25 amps. The calibrating circuit range gives voltage values across the 10 ohm calibrating resistor (Rc1) from approximately .04 to .14 volts, equivalent to currents from 4 amps to 14 amps across a .01 resistor. This was perfect for my application, which was set for 10 amps (my load normally draws 8 amps). Any one of the three calibrating resistors may be changed to give you the value you need.

Calibrating the circuit is quite simple. Connect the three calibration resistors and adjust them so there is a voltage drop across Rc1 equal to 1/100th of the amount of amps you wish to set the circuit breaker for, i.e., use .01 volts for a 10 amp trip point. Connect Rc1 where Rsense will be connected later. Now adjust R15 (25K 10 turn Trimpot) just to the point where the LED turns on. Now push the reset button, which will turn the LED off, and adjust R15 again until the LED turns on. Remove Rc1, push the reset button, and reconnect Rc1. If all has gone well, the LED will turn on again. You are now within plus or minus .5 amps of the 10 amp trip point. Since the project is only meant as a circuit breaker, not a

Figure 4. Circuit Breaker Schematic.

Parts List

Resistors	Source	Part Number	IC		
R7 = 1K	Digi-Key	1.0KEBK-ND	LM3900 Quad Norton op-amp, 14 Pin DIP	Digi-Key	296-9565-5-ND
R1,R2,R8 = 10K	Digi-Key	10KEBK-ND			
R13 = 33K	Digi-Key	33KEBK-ND			
R5 = 82K	Digi-Key	82KEBK-ND	**Miscellaneous**		
R4,R9 = 100K	Digi-Key	100KEBK-ND	NO PB = Normally Open Momentary Push Button Switch	Digi-Key	506PB-ND
R11 = 160K	Digi-Key	160KEBK-ND			
R12,R14 = 560K	Digi-Key	560KEBK-ND			
R6 = 1M	Digi-Key	1.0MEBK-ND	LED, Panel Mount	Digi-Key	67-1171-ND
R10 = 2.7M	Digi-Key	2.7MEBK-ND			
R3 = 3.9M	Digi-Key	3.9MEBK-ND	14 Pin DIP Socket (for LM3900)	Digi-Key	ED58143-ND
R15 = 25K Trimpot	Digi-Key	490-2269-ND			
Capacitors			Single Sided copper clad PC Board, 3"x 4"	Digi-Key / RadioShack	PC9-N9 / 276-1499
C3 = .033mf	Digi-Key	P4954-ND			
C5 = .1mf	Digi-Key	P4924-ND			
C1,C2,C4,C6 = .22mf	Digi-Key	P4966-ND	PC Board (similar to Vector Board) 356 Holes, 1 7/8 inch x 2 7/8 inch	RadioShack	276-149
Transistor					
2N3904 NPN or Equivalent	Digi-Key	2N3904-D26ZCT-ND			

current limiter, you will probably want to adjust it on the high side so that it takes a bit more current to trip the breaker.

I tried to purchase a .01 resistor capable of up to 150 watts or so, but could not find one. I decided I could proba-

bly use the sort of steel wire that is normally used for putting up wire fences. I was able to find 14-gauge wire at Home Depot, which turned out to have a resistance of .0009 ohms per centimeter. Discovering this was an interesting exercise using Ohm's Law. I used a 12 volt car headlight as a load, which used about 8 amps from a 12 volt auto battery. By putting a 10 cm length of this wire in series with the headlight/battery circuit and measuring the millivolt drop across the wire, I was able to determine the resistance of the wire. Fabricating a .01 resistor was pretty easy after that, since .0009/.01 = 11 centimeters. I made the wire into a single 1 cm diameter winding, and I had my Rsense!

I constructed a printed circuit board for this project (see Figures 5 and 6 on the *Nuts & Volts* website at **www.nuts volts.com**), but it could also be breadboarded using vector board. I remotely located the FET and Rsense so that I could mount it to my chassis as a heatsink (which was probably unnecessary, since the FET never even gets warm). Using a printer to print out the circuit board works nicely, giving the proper size of the board if you print it using the original DPI size. The result may easily be used for making the PC board.

There is an excellent viewing and printing program available for free from **www.irfanview.com** When printing with this program, it gives the option to print using the original DPI of the file. Note that "Japanese" type construction was used where most of the parts are mounted vertically to give a smaller PC board size. This method accommodates different size parts much more easily than horizontal mounting.

Note that the data sheets for the LM2900/3900 and IRF1404 FET, as well as any circuit board layouts will be posted on the *Nuts & Volts* website.

All of the parts for the project are available from Digi-Key (**www.digikey.com**) except for the vector board. Many of these parts are also available from RadioShack. **NV**

What's the difference?

Price!

Build a Basic MP3 Player Platform: Part 1

Design Allows for Easy Customization

Player With Display in Plastic Case.

Ever have the desire to build your own embedded MP3 player? Maybe you just like to listen to music and want to do it with a gizmo you built yourself, or maybe you don't like the user interface on your commercial MP3 player and are sure you could design a better one. Or maybe you want the ultimate recorded voice for your robot project, or how about a fancy doorbell that plays an entire song? Or record an assortment of different dogs barking and connect it to your burglar alarm?

Not too long, ago this would have been a major undertaking, but now the technology has evolved to the point where any hobbyist can easily put together an MP3 player. Here's a basic MP3 player platform with a design that's intentionally left open to allow customization by the builder.

It can play MP3 files from a CompactFlash™ card at any bit rate up to 320 kbps. It supports either a LCD or vacuum fluorescent display for menus and MP3 information, and a rotary encoder plus a push button for user input. The microprocessor used is a Philips 8051 derivative with Flash memory that can be programmed in system using a standard RS-232 port; no special programmer is needed.

The firmware is open source and is released under the GNU Public License so that you can customize it for your own needs.

Although it's easy enough to assemble, this is not a trivial project, so we'll present it in two parts. In the first part, we'll show the schematics and discuss the operation of the hardware and software. In Part 2, we'll cover construction and assembly details, checkout and testing procedures, give instructions for operating it, and discuss

Figure 1. Audio Schematic.

a few possibilities for modification.

Hardware

The heart of the hardware is the ST Microsystems STA013 MPEG Audio Decoder chip. After it's properly configured by the software, this chip is able to accept any MP3 file and decode it into a serial I2S digital audio stream. The CS4334 serial audio DAC takes the digital audio bit stream and converts it to analog, which is then filtered and output via J2. Together these two chips can decode any MP3 file, fixed or variable bit rate, up to 320 kbps and 96 kHz sampling rate. Although this is a basic player design, the resulting audio quality is as good as any.

The STA013

Figure 1 shows the portion of the circuit associated with the STA013 and CS4334. The RC network at the output of the CS4334 — R6 through R9 and C8/C9 and C20/C21 — are a simple low pass filter that limits the output to audio frequencies. Note that this part of the circuit has a specially filtered power supply — L1 and C10/C11 — as well as a separate analog ground, in an attempt to keep digital switching noise out of the audio output. By the way, this part of the circuit is straight out of the Cirrus Logic data sheet for the CS4334.

The STA013 requires a 14.31818 MHz clock for timing. In principle, the STA013 can connect to a crystal directly, but in my prototypes I had difficulty getting the STA013's internal oscillator to start reliably, so I used a packaged TTL oscillator, OSC1, instead. The STA013 actually uses an internal PLL (Phase Locked Loop) locked to this external crystal to generate all its internal timing; components C6, R5, and C7 are a loop stabilization filter for this PLL and R1, R3, and C3 filter the PLL power supply to prevent noise from causing glitches in the timing. These parts are all taken directly from the ST Microsystems STA013 application note, AN1090.

Interfacing 3.3V and 5V Logic

The problem with the STA013 is that it is strictly a 3.3 volt part; that's an issue because the rest of our MP3 player runs from a 5V supply and it makes interfacing the STA013 inputs and outputs with the rest of the circuit a tad tricky. STA013 outputs that drive 5V inputs — e.g. the SDO, SKCT, LRCKT and GCLK pins that drive the CS4334 — can simply be connected directly. When the STA013 output is low it will be at nearly 0V, which is a logic low anywhere, and when the STA013 output is high it will be almost 3.3V. The standard CMOS threshold for a logic high is 1/2 VCC, or 2.5V for the CS4334, so 3.3V will be seen as a high level. The noise margin, 3.3V-2.5V, is not as good as you would get from a 5V output, but it's still adequate.

The same is true of the STA013 DATA_REQ output which drives a microprocessor input pin, but in this case a 4.7K (R17) series resistor is used to limit the current that would flow if the software accidentally programs this microprocessor pin as a high level output. Applying 5V directly to a STA013 pin would destroy it, and in the case of microprocessor outputs that drive STA013 inputs, such as RESET, we use another 4.7K resistor (R11) to limit the current that flows and the voltage on the STA013 pin to safe levels.

In principle, the STA013 SDI and SCKR inputs present the same problem as RESET, but these are high speed signals and we bypass the series resistor with a small capacitor (C12-13) to improve the rise time. The last two pins for the I2C bus, SDA and SCL, are a special problem. I2C is a bidirectional bus with open drain drivers and requires pull up resistors to VCC. A general-purpose voltage shifter for I2C is very difficult, but in this case we can cut a few corners because there are only two chips involved — the STA013 and the microprocessor — and the microprocessor is always the bus master.

Since the bus master (the microprocessor) always drives the clock, we can pull up SCK to 5V and use a simple series resistor (R13) like RESET/SDI/SCKR. SDA is harder to handle because it is completely bidirectional —

the simplest thing to do is to pull it up to 3.3V; this won't hurt the STA013 and is enough to make the microprocessor happy. R12 limits the fault current that flows should the microprocessor pin be accidentally taken to 5V — with I2C this never happens, but a software bug could wrongly change the port pin programming. Good hardware design dictates that bad software should never be able to physically fry chips!

The Microprocessor

The microprocessor used is a Philips 89C664 with 16K of Flash programmable memory and 4K of internal SRAM. An 89C668 can also be used; the only difference being that this part has more memory. This MCU was selected because a) it has a fairly large amount of on chip RAM, which is needed for buffering, and b) because it can be programmed and reprogrammed "in circuit" by downloading the firmware directly from a PC over a serial RS232 port. No special programmer is required. This makes it easy to modify and update the firmware, which I hope will encourage people to enhance the

player's firmware.

Figure 2 shows the microprocessor and associated circuitry. Everything is very straightforward, except for the programming port — U2, J7, R10, and Q1. The programming cable you make, which connects J7 to the PC's serial port, should connect J7 pin 4 to ground (J7 pin 1). This enables Q1 to pull the PSEN pin low while reset is asserted, which is the 89C664's signal to enter programming mode. After you've programmed the microprocessor (more on that in Part 2), you must disconnect the programming cable to disable Q1 and allow the PSEN pin to float. If you don't do this, then the microprocessor will stay in programming mode and won't run the firmware you've downloaded!

Transferring MP3 Data

The microprocessor reads the data from the MP3 file on the CompactFlash card and transfers it to the STA013. The STA013 wants this data serially, one bit at a time; the SDI input is the serial data input and SCKR is the serial clock. JP2 and JP3 (Figure 1) allow these

Figure 2. CPU Schematic.

pins to be connected to either the 89C664's port 1, bits 1 and 2 (P11 and P12), or to the 89C664's serial port RxD and TxD pins. In the latter case, JP1 should be removed to disconnect the PC's serial data from the microprocessor.

Why would you want to connect the STA013's serial data input to the microprocessor's UART? It's because all 8051 chips, including the 89C664, have a mode which allows the internal UART to be used as a simple shift register. In this mode, the TxD pin becomes the clock output and RxD is the data output. The advantage to doing this is that the 8051's hardware UART is vastly faster then generating SDI and SCKR in software. Remember, when we're playing 320 kbps MP3 files, we have to send the STA013 at least 320K bits (that's 40K bytes!) every second — that's more than the software can manage by "bit banging" the P11/SCKR and P12/SDI pins.

The drawback is that we also need the 89C664's UART for downloading software and for connecting to a regular terminal when we're trying to debug the firmware. Jumpers JP1-3 allow you to choose, and by changing a compile time option you can rebuild the firmware to output MP3 data with P11 and P12 or the internal UART.

I/O Interfaces

Figure 3 shows the remaining I/O connectors on the MP3 player board. J9 is a 50 pin CompactFlash card socket; grounding pin 9, OE, of the CF card forces it into TrueIDE mode, which emulates a standard IDE disk drive. However, unlike most modern IDE drives, all CF cards are required by the standard to support an eight bit interface, which simplifies connecting the CF card to the micr processor.

J5 is for connecting an eight bit parallel alphanumeric vacuum fluorescent display (VFD) and J4 is used for connecting an alphanumeric LCD display. Of course, only one of these two should be used at any one time! The pinout of the 14 pin LCD connector is pretty standard and any display that uses the Hitachi 44780 controller chip should work. Unfortunately, VFD displays are not as standardized — this pinout works with many Noritake and Futaba displays, but you should double-check the pinout of your display before connecting it.

J6 is for connecting an optical rotary encoder with a built-in push button. The rotary encoder specified has a wonderfully smooth feel but can be rather expensive. You can substitute a cheaper, mechanical encoder if you want.

Figure 3. I/O Connectors Schematic.

Figure 4. Power Supply Schematic.

Power Supply

Finally, Figure 4 shows the power supply section. This generates both 5V and 3.3V from a 9-15V unregu-lated input. D2 protects against reverse input polarity and D3 protects against damaging voltage spikes; that's especially important if you intend to use your MP3 player in an automobile with their notoriously noisy electrical systems.

J4 can be used to supply power to a VFD display, but if you do this be aware that most VFDs require 300 to 500 mA, which will increase the heat dissipation of VR2 considerably. Without a VFD display, VR2 will barely get warm and no heatsink is required. With a VFD display, VR2 will get quite hot and will require the biggest heatsink you can manage. LCD displays use very little current and have no effect on VR2.

Software

The software consists of some 7,000 lines of C and assembler code in about 20 different source files. It is able to access the CompactFlash card using TrueIDE mode, read a FAT16 file system, and locate and read all the MP3 files in the root directory. It also configures the STA013, sends it MP3 data from the file, shows the current file and ID3 tag information on the attached display, and responds to user input through the rotary encoder. Finally, the software contains a basic power on self test (a.k.a. POST) which tests the hardware for common failures and defects every time power is applied.

It's a sad truth of modern electronics that the software is usually more difficult and far more time consuming than the hardware, and this project is no exception. The firmware represents pretty much the absolute minimum amount of functionality that's required to have anything useful. On the happy side, though, I have made the firmware open source and released it under the GNU Public License. Interested readers are encouraged to enhance it and fix any of the limitations that prove unbearable.

FAT.C

This module contains the code needed to read a Microsoft™ FAT16 file system. It supports only a very simplified set of FAT functions and, even at that, it's already the longest single module in the firmware. It's simplified because:

• It only supports reading (an MP3 player has no need to write files, after all!).

• It only supports FAT16 (modern CF cards are too large for FAT12 and FAT32 is much more difficult to implement).

• It doesn't support a master boot record (i.e., the CF card can't be partitioned).

• It doesn't support subdirectories (all MP3 files must be kept in the root directory).

BUFFER.C

The internal RAM of the 89C664 is initially divided up into a pool of 512 byte buffers. The actual "player" part of the firmware is structured as two independent asynchronous processes. The background process pulls empty buffers from the free pool and fills them with data from the MP3 file. It blocks when there are no free buffers, and it stops when we reach the end of the file.

The foreground process, driven by the STA013 data request interrupt, pulls full buffers from the "play" list and sends them serially to the STA013. Hopefully, it never finds the play list empty before the file ends; the audio will glitch if the STA013 runs out of data. After the data is sent to the STA013, buffers are returned to the free pool.

MP3DRQ.A51

This module sends the MP3 data serially to the STA013 chip. It can be configured, by a conditional assembly switch, to send the data using either the 89C664's internal UART or by "bit banging" port bits P11 and P12.

I2C.A51

This is a simple software I^2C implementation that's used to send commands to the STA013. It may seem foolish to use a software I^2C implementation when the 89C664 has a perfectly good internal hardware I^2C controller, but the firmware was originally written with a cheaper microprocessor in mind. None of those proved to have enough internal RAM for buffering and so we ended up with the 89C664 and an unused I^2C controller. You're welcome to fix this, however, since the STA013 uses the

I²C only for simple commands (e.g., STOP, PLAY, etc.) and not for the actual MP3 data; it has little effect on the performance.

PLAYER.C

This module is the "main program" for the firmware. On power-up, it configures the hardware, executes the post, reads the directory from the CompactFlash card, plays MP3 files, and responds to the rotary encoder and button user interface. All of that, of course, is done with the help of the other modules!

Other Modules

Here's a brief list of some of the remaining software modules:

• CFCARD.C — CompactFlash Flash memory card interface functions.

• IDE.A51 — Low level (identify device and read sector) IDE functions for the CF card.

• DISPLAY.C — Drivers for the VFD and LCD displays.

• POST.C — Power-on self test functions.

• STA013.C — STA013 initialization, configuration, and control functions.

• SWITCHES.C — Rotary encoder and push button interface.

• TIMER.C — Time keeping functions.

• DEBUG.A51 — Sends debugging messages to the serial port.

Until Next Time

That's all for this issue. Next month, we'll actually build and test our MP3 player. In the meantime, you can visit the Spare Time Gizmos MP3 Player webpage — http://mp3.SpareTimeGizmos.com — to read up on the latest news of the project. **NV**

A Whimsical Doorbell

Combining Folk Art and Wireless RF Technology

The use of radio frequency (RF) control for switching devices is becoming increasingly popular. The project described here uses RF control to actuate a zany, whimsical doorbell. The electronics involved in this circuit can be used to control more practical projects, such as a garage door opener, appliance turn-on, wireless home signaling (paging), and remote sensing.

Since my knowledge of RF control was limited, I was delighted to discover, through ads in *Nuts & Volts* and on the Internet, that ready-made RF transmitter and receiver modules were readily available. These would perform exactly what was needed at low power and low cost.

There are two parts to this project. One part is the doorbell actuator, where one actuates the transmitter circuit and sends out an encoded RF signal. In this unit, the actuation device is a woodenhead carving of a man with his tongue sticking out (Figure 1). Pressing the tongue will actuate the transmitter circuit. Most people are a bit repelled by that. But that is the purpose of folk art — to do something unique and personal. The transmitter is located in the actuation device and is powered by a nine-volt battery. Because the actuator is battery-operated, it is free to be placed anywhere. It was tested to a distance of 300 feet.

The second part of the project is the actual doorbell unit as described above. It consists of four bells mounted on a strap and an additional bell on the top of the façade, pealed by a Christmas elf (Figure 2).

The RF receiver is located in the doorbell unit and it is powered by a 24-volt wall wart. The reason for his is that the gear motor used to shake the bells came from an old copying machine and required 24 volts to operate.

One requirement was some sort of indication of battery voltage level, since a doorbell would be of no use if the actuator battery was dead. So, included in the transmitter circuit is a circuit that detects battery voltage. The circuit lights a green LED in one eye of the head carving when the voltage is within range and a red LED in the other eye when the voltage is below range. These LEDs go on only when the unit is activated.

Transmitter and Receiver

I used the Reynolds Electronics (**www.rentron.com**) TWS-434A transmitter ($8.50) and RWS-434 receiver ($8.50) (see Figure 3). These are inexpensive devices that, along with their associated encoder (HT-12E, $1.90) and decoder (HT-12D, $1.90), allow one to transmit and receive a control signal. The transmitter can control one out of 256 devices with a "turn on" signal. The receiver must be set to the same channel that the transmitter transmits on. This is done by the use of an eight-bit code bank switch on both the transmitter and receiver circuits.

The transmitter and receiver pair is used to send a "turn on bell" signal. Nothing fancy, just set the bell ringing. Releasing the pushbutton or, in this project, releasing the tongue, stops the RF signal and switches off the bells.

At first, the channel was set to "0." Occasionally, at random times, the bell would ring spontaneously. It was picking up the signal from my home weather station unit that had several remote senders. When the channel was changed to 13, the bells worked perfectly and the receiver never sent another errant signal to the bells.

Figure 1. Carved head actuator. Press tongue to operate.

Figure 2. Angelmaker exterior view. The Christmas elf can just be seen at the top left.

The Transmitter Circuit

The Reynolds Electronics TWS-434A transmitter module transmits at 434 MHz (see Figure 4). The transmitter must be used with the HT-12E four-bit encoder IC. This encoder, upon receipt of a transmission enable, will begin a four-word transmission cycle. This cycle will repeat itself as long as the transmission enable is held low. The eight-switch code switch band chooses the channel desired. As previously mentioned, the receiver code switch bank must be set to the same configuration in order to receive the signal.

An encoder is necessary because the receiver is always on and it detects any and all signals that enter the antenna. Much of it is random noise. To separate the noise from the true signal, some sort of encoding must be performed. This is done by the HT-12E chip.

The pushbutton terminals on the printed circuit board are connected to a miniature snap switch that is actuated by pressing down the tongue on the carved head (see Figure 1 again). The switch connects the nine-volt battery to the voltage regulator, IC-1, which sends a regulated +5 volts to the RF transmitter module and

Figure 3. Transmitter and receiver printed circuit boards. Notice that both code switches are set to the same code — 10110001 when read from left to right. Since B0 on the PC board is on the left, the code really is: 10001101, or 141 decimal.

encoder. Since the TX pin on the encoder is permanently wired to ground, an encoded RF signal is sent out of the antenna whenever the transmitter is powered up. A simple nine-inch 22-gauge insulated wire is used as the antenna.

Figure 4. Transmitter schematic.

Figure 5. Receiver schematic.

Figure 6. Rear view of carved head actuator. The prototype receiver circuit board is shown at the top. Below the board is the battery and below that is the tongue actuator and snap switch.

Battery Level Check Circuit

When using a battery-operated transmitter to power a doorbell, it would be convenient to know if the battery is too low to operate the transmitter. So, a battery indicator circuit was designed. This is performed by IC-3, the LM339 Quad op-amps, wired as voltage comparators. Only two of the op-amps are used for this circuit function.

The green LED lights every time the pushbutton is depressed when the nine-volt battery is above 6.8 volts. The red and green LEDs light when the battery voltage goes below 6.8 volts and above 6.3 volts. But when the battery goes below 6.3 volts, the red LED lights and it is time to change the battery. I estimate that the nine-volt battery will last about two years or nearly shelf life, that is, unless you get an inordinate number of visitors.

The Receiver Circuit

The Reynolds Electronics RWS-434 Receiver Module picks up this signal and the HT-12D four-bit decoder compares the serial input data three times continuous-

Figure 8. View of the carved head actuator.

Figure 7. The Angelmaker interior view. The operating mechanism and the prototype receiving circuit can be seen here. The gold and green cam shakes the bells and activates the Christmas elf using the steel wire seen going up to the right. The antenna wire is the horizontal blue wire near the top of the green back board. The gold cam is powered by the 24-volt copy machine motor.

ly (see Figure 5). If no error or unmatched codes are found, the input data codes are decoded and then transferred to the output pin, VT. This eliminates any spurious turn-on by random RF noise signals.

The VT output then drives transistors U1 and U2. U1 is used as a buffer to drive a "Receive Signal" LED. The Darlington transistor, U2, is used to drive a load of up to 4 amps at 9 to 30 volts, depending upon the power supply used. The output of the decoder is a positive logic signal and, by using a buffer transistor, one could connect virtually anything to it. This could be a solid-state relay for 120-volt control, a reed relay, LEDs, and so forth.

Notice that there are two positive power input terminals. One is for a nine-volt system and the other is for a higher voltage of 24 to 30 volts. Diode, D1, drops the voltage so that the voltage regulator, IC-1, does not have to dissipate so much power.

Antenna

As can be seen in Figures 6 and 7, the antenna for the transmitter and receiver, respectively, is simply a nine-inch length of solid insulated wire running around the head in the actuator and around the frame in the bell unit. The unit

Transmitter Board Parts List

Resistors (1/4 watt, 5%)
R1 750K
R2 2K
R3 1K
R4 220
R5 10K

Semiconductors
D1, D2 1N4148 signal diode, Jameco #36038
D3 Zener, 5.1V, 1/4 watt, Jameco #259493
IC-1 LM78L05 Voltage Regulator, 5V TO-92 case style, Jameco #51182
IC-2 Four-bit Encoder, Reynolds Electronics HT-12E
IC-3 LM339, Jameco #143888
IC-4 Transmitter module, Reynolds Electronics #TWS-434A

LED1 Green
LED2 Red

Miscellaneous Items
SWB-1 Dip Switch, eight poles, Jameco #38842
PCB-1
RF transmitter board
Nine-volt battery connector, Jameco #216451
Antenna 9.4" 22-gauge solid wire

Project

Receiver Board Parts List

Resistors (1/4 watt, 5%)
R1   33K
R2   220
R3   10K
R4   1K

Capacitors (Aluminum Radial-lead +/- 20%)
C1   1 uF   50V   Jameco #94211
C2   10 uF   25V   Jameco #29832
C3, 4   .01 uF   50V   Jameco #15230

Semiconductors
D1   NTE5127A zener diode, 12V, 5 watts, Jameco #274335
D2   1N4004 400 PRV, 1 amp, Jameco #35991

LED 1   Green

U1   2N3904 NPN, Jameco #38359
U2   TIP 120 or 121 Darlington, Jameco #32993

IC-1   LM78L05 voltage regulator, 5V, TO-92 case style, Jameco #51182
IC-2   Receiver module, Reynolds Electronics RWS-434
IC-3   Decoder, Reynolds Electronics HT-12D

Miscellaneous Items
SWB-1   Dip switch, eight poles, Jameco #38842
PCB-1   Etched printed circuit board (RF receiver)
Power source   (depends on application. 9V to 30V wall wart)
Antenna   9.4" 22-gauge solid wire
4-40 x 1/4" long machine screw, lock washer, and nut for mounting T2

was tested at 300 feet away. This gives plenty of range for most needs.

Conclusion

It is not difficult to use wireless control for your own projects. Another planned application is to control the existing driveway and outdoor house lights from my automobile. All that would be required is to design a simple latch and delay circuit that is activated by the VT signal of the decoder. Instead of the Darlington driver, I would use a solid-state relay.

The art project described here is the first of what I hope will be a series of comical doorbells. Some projects will involve the playing of tuned bells and others will have tuned organ pipes or slide whistles. An earlier piece was recently installed in a folk art gallery in north Georgia. Using RF wireless control has opened many new and exciting doors for artists and experimenters. **NV**

This is where we live!

Photonics Rules of Thumb
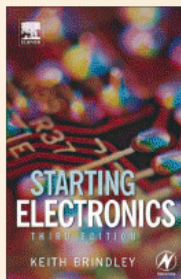### by John Lester Miller / Ed Friedman

Even the brightest and most capable optical engineer would be stymied without the use of practical principles of estimation. These ``rules of thumb" make it possible to get a quick idea — completely adequate for the moment — of the impact a given design change will have. Now, for the first time anywhere, these time-tested estimation rules have been gathered together in a single handy source. Arranged alphabetically according to specialty, this unique book spans the entire spectrum of photonics, from optics to lasers. Scientists and engineers at all levels will want to keep this fast and easy-to-use reference near at hand. **$59.95**

Starting Electronics: Third Edition
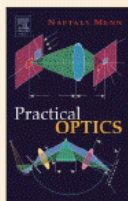### by Keith Brindley

Starting Electronics is unrivalled as a highly practical introduction for hobbyists, students, and technicians. Keith Brindley introduces readers to the functions of the main component types, their uses, and the basic principles of building and designing electronic circuits. Breadboard layouts make this very much a ready-to-run book for the experimenter; and the use of multimeter, but not oscilloscopes, puts this practical exploration of electronics within reach of every home enthusiast's pocket. New material includes sections on transducers and more practical examples of digital ICs. **$19.95**
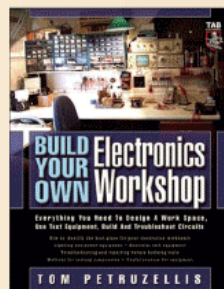
Practical Optics
by Naftaly Menn

Practical Optics bridges the gap between what students learn in optical engineering curriculum and the problems they realistically will face as optical engineering practitioners, in a variety of high-tech industries. A broad range of topics from modern optical engineering are addressed, ranging from geometrical optics and basic layout of electro-optical instruments, to the theory of imaging, radiation sources and radiation sensors, spectroscopy, and color vision and description of optical methods for measurements of numerous physical parameters (temperature, flow velocity, distances, roughness of surfaces, etc.). **$69.95**

Build Your Own Electronics Workshop
by Thomas Petruzzellis

The Electronics Workbench was created to assist the newcomer in the field of practical electronics through the creation of a personal electronics workbench. It is a place specially designed so that readers can go there to work on an electronic project, such as testing components, troubleshooting a device, or building a new project. The book includes invaluable information, such as whether to buy or build test equipment, how to solder, how to make circuit boards, how to begin to troubleshoot, how to test components and systems, and how to build your own test equipment, complete with appendix & resources, etc. This is *THE* book for anyone entering the field or hobby of electronics. **$29.95**
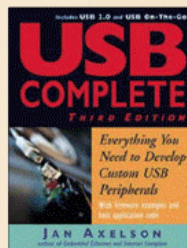
Home Computing
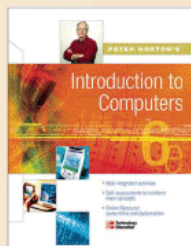
USB Complete, Third Edition
by Jan Axelson

Every recent PC has Universal Serial Bus (USB) ports. In *USB Complete*, Jan Axelson shows how to design and program devices that use USB to communicate with PCs. Learn how to select a USB controller chip that fits your project and budget, how to write program code to manage USB communications in your device, and how to write Visual Basic .NET and Visual C++ .NET applications that exchange data with your devices. Explore how the interface works from the ground up and learn how to create devices that communicate reliably with all PCs. **$49.95**

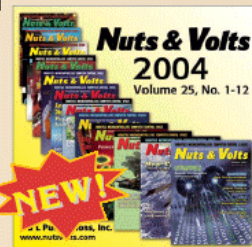## Peter Norton's Introduction to Computers 6/e
by Peter Norton

This completely revised edition of *Introduction to Computers* offers a comprehensive grounding in computer concepts, reflecting on some of the latest advances in technology, to prepare students for the next level of computer productivity. It includes three entirely new chapters on the Internet, computer graphics, and multimedia. **$58.75**
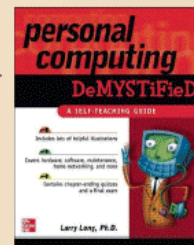
Nuts & Volts CD-Rom

Here's some good news for *Nuts & Volts* readers! Starting with the January 2004 issue of *Nuts & Volts*, all of the issues through the 2004 calendar year are now available on a CD that can be searched, printed, and easily stored. This CD includes all of Volume 25, issues 1-12, for a total of 12 issues. The CD-Rom is PC and Mac compatible. It requires Adobe Acrobat Reader version 6 or above. Adobe Acrobat Reader version 7 is included on the disc. **$29.95**

Personal Computing Demystified
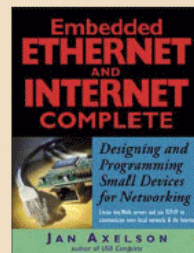### by Larry Long

Quickly learn how to get the most out of your personal computing experience, whether you want to network multiple PCs, buy a PC, store photos, or download music. This easy-to-understand, step-by-step guide walks you through each task you need to complete without using unnecessary technical language. A user-friendly resource, you don't need to have previous experience to get fast results. **$19.95**

Embedded Ethernet and Internet Complete
by Jan Axelson

Learn how to design and program devices that host Web pages, send and receive e-mail, and exchange files using FTP. Put your devices on the Internet and monitor and control your devices from across town or around the world. Or create private, local networks that enable your devices to share information, send commands, and receive alarms and status reports. Plus: learn about Ethernet controllers, hardware options for networks, networking protocols, and how to keep your device firmware and data secure. **$49.95**

RELOCATION ASSISTANCE

by Al Williams WD5GNR

If you program PICs in assembly language, it is a good bet you use MPASM — Microchip's free assembler/editor/simulator. It's also a good bit that you write one big file that maps out the program in the PIC's memory.

The single file approach is great for small programs. Besides, when you started out with the old 16C84, all programs were small — the chip only has 1K of program space! But if you've moved to newer parts like the PIC16F877, you have to fill 8K of

program space so your programs can get a bit more complicated.

Luckily, Microchip gives you an option. MPASM can generate relocatable object files. Another program, MPLINK, can combine multiple object files and create a single hex image you can download to your PIC. MPLAB (Microchip's development environment) can control the whole process for you.

This scheme has several advantages. First, it allows you to partition your program into separate parts which is a great benefit when your program is large. Even more importantly, with a little planning, you can create stand-alone modules that are reusable. Then you can easily link a pretested module in with many different programs, saving time, effort, and debugging headaches.

Every PC programming book starts with a Hello World program. For microcontrollers, the equivalent is a blinking LED program. To that end, I'll show you how to make a flashing LED program using relocatable modules on a PIC16F873. The project is set up to work with the Microchip ICD (see the November 2001 issue), an APP-II (**www.awce.com/awce/app2. htm**), or a plain PIC16F873. In any event, you just need an LED (and the normal dropping resistor) connected to Port B pin 0.

Getting Started

The first step is to create a new MPLAB project (using the Project Wizard on the Project menu). Make the following choices:

- 16F873A (or whatever processor you are using)
- Toolsuite: MPASM
- Project name and directory as you wish (I used objtest)
- Add the 16F873A.LKR (or equivalent) file from the MPASM Suite\LKR directory (normally under c:\Program Files\Microchip (you will want to copy this file to your project file so you can modify it and control it)

When MPLAB notes that you have an LKR file, it uses relocatable mode. Instead of directly building a hex file from an assembly language file, the IDE builds object files from each assembly file and then uses the linker (MPLINK) to join them together into a hex file. The LKR file describes the memory areas available on the chip.

If you already have ASM files, you can add them via the Project Wizard, or add them later by using the Project | Add Files to Project menu item.

You can create completely custom versions of the LKR file, but usually it is better to use the predefined ones that Microchip supplies (at least until you get comfortable with relocating code). If you make a private copy

---

LISTING 1. An LED flasher that uses relocation.

```
LIST P=16F873A
#include "P16F873A.INC"
#include "delay.inc"
#include "relocate.inc"



STARTUP  Code 0x0
         nop       ; make ICD happy
         PAGESEL Start
         goto Start

         Code
Start:   movlw 0xF0
         BANKSEL delayct
         movwf delayct+1
         movlw 0
         movwf delayct ; delay=0xF000

blink:
         BANKSEL TRISB
         bcf TRISB,0   ; make RB.0 an output
         BANKSEL PORTB
loop:
         bsf PORTB,0
         xcall Delay
         bcf PORTB,0
         xcall Delay
         goto loop

         END
```

of the file, you can at least change it if necessary.

There are scripts in the MPASM Suite directory for all the processors. Sometimes there will be special scripts for different circumstances. For example, the 16f873a.lkr script is for a normal 16F873A, while the 16f873ai.lkr script is for the same chip when using the ICD or anything that reserves the same resources as the ICD (like the APP-II).

In fact, I didn't agree with the shared banks defined in the standard linker file, so I changed them (see below for more details). You may want to make other changes to maintain control over memory allocation.

With your project set up properly, MPLAB works like it always does. You build your code and you can simulate it, download it (if you have a PICStart programmer), or simulate it on an ICD2 or other MPLAB-compatible simulator. The problem is: What goes in the ASM files?

Inside the Relocatables

When you are generating relocatable code, you follow different rules than you are used to from working with MPASM in single-file mode. Since each module has to coexist with other unknown modules, you can't expressly reference code and data locations anymore. Instead, you use symbolic labels. That means no more ORG directive! A relocatable ASM file consists of multiple sections.

One of the most common sections is a CODE section. As you'd expect, your program's code is in there. The CODE section has a label which is the section's name (if you don't provide a name, the assembler uses .code as the default). You can put multiple sections in one file as long as they have different names. The linker will put all the sections with the same name together. Also, the linker script will define sections (for example, page0 and page1). If you use one of these names, it will indicate that you want that code section in that physical address range. You might write:

```
CODE
Main:   clrwdt
```

Or:

```
page0   CODE
```

Optionally, you can also specify an address for the CODE section. This is useful for specifying interrupt vectors or other items that must be a fixed address. For example:

```
RESET CODE 0x0
```

In addition to code sections, your files can create data sections for variables. The UDATA section contains unitialized data — the most common type. The linker will assign each variable unique space in the PIC's registers. For example, consider this code:

```
        UDATA
aByte   res 1
aWord res 2
```

This sets up a byte variable and a 16-bit variable. The exact addresses aren't known until link time, but your program can still use these symbols in instructions. In fact, you can even add constant offsets to the symbols (aWord+1, for example) or use the HIGH or LOW prefixes to select one byte of a 16-bit word. You can also subtract two symbols that are in the same section (aWord-aByte, for example). You can specify initialized data using IDATA instead of UDATA, but this isn't very useful. PICs start up with random garbage in their memory, so you have to write your own code to initialize IDATA variables. That means there isn't much benefit to using them — just define your variables in the UDATA section and initialize them in your program prior to using them.

Another useful data section is the UDATA_SHR section. Variables defined here will be in memory that is visible from all banks. The UDATA_OVR section contains variables that you don't mind other modules writing over. This is useful for temporary variables that you don't need for very long.

Note that many of the Microchip-supplied LKR files specify unshared memory as PROTECTED. This means you must specifically name that section in order to use it. Also note that UDATA expects DATABANK keywords in the LKR file. UDATA_SHR expects SHAREBANK keywords.

Intermodule Communications

In general, modules can't see symbols defined in other modules. Of course, you want some symbols to be visible, so you must declare them using the GLOBAL keyword after you define them. For example:

```
        UDATA
ctr     res 1
limit   res 1
        GLOBAL ctr, limit

        CODE
DoCount
        GLOBAL DoCount
```

In addition, the module that wants to use these symbols must declare them as EXTERN, like this:

```
        EXTERN ctr, limit, DoCount
```

Banking

There is one problem with not knowing symbol values at compile time. What do you do about code and data banks? The PIC uses multiple memory banks to cut down on the size of the addresses required. So a jump to location 0x10 might really be a jump to location 0x10 in any of two or four pages depending on the processor.

To solve this problem, the assembler allows you to write PAGESEL to set the processor's page bits for a particular code symbol and BANKSEL to set the data bank to the right value for a particular data symbol.

Just like regular PIC programming, this requires a bit of discipline. I usually force the page bits on each call and then reset them to the current page so that any unadorned CALLs or GOTOs will stay on the same page. MPASM even includes a pseudo instruction LCALL (and the corresponding LGOTO) that is equivalent to a PAGESEL and a CALL (or GOTO).

An Example

My example program appears in Listings 1 and 2. Listing 1 is the main program. It sets up the reset vector and the I/O ports. The main program includes delay.inc (Listing 3) which imports the global's from delay.asm. The modified linker script appears in Listing 4. I removed the protected attribute from the SFR banks so that MPASM and MPLINK would automatically populate the UDATA sections.

The include file strategy is similar to how you do includes in C and it makes it very easy to add a module to another program later down the road. The main program includes the relocate.inc file (Listing 5) which has a macro, XCALL, that is useful for calling across banks.

> **LISTING 3.** The header file for the Delay routine.

```
; include for delay.asm
        EXTERN Delay, delayct
```

> **LISTING 4.** The modified linker script.

```
// Sample ICD2 linker command file for 16F873a

LIBPATH   .

CODEPAGE    NAME=vectors    START=0x0       END=0x4         PROTECTED
CODEPAGE    NAME=page0      START=0x5       END=0x7FF
CODEPAGE    NAME=page1      START=0x800     END=0xEFF
CODEPAGE    NAME=debug      START=0xF00     END=0xFFF       PROTECTED
CODEPAGE    NAME=.idlocs    START=0x2000    END=0x2003      PROTECTED
CODEPAGE    NAME=.config    START=0x2007    END=0x2007      PROTECTED
CODEPAGE    NAME=eedata     START=0x2100    END=0x217F      PROTECTED

DATABANK    NAME=sfr0       START=0x0       END=0x1F        PROTECTED
DATABANK    NAME=sfr1       START=0x80      END=0x9F        PROTECTED
DATABANK    NAME=sfr2       START=0x100     END=0x10F       PROTECTED
DATABANK    NAME=sfr3       START=0x180     END=0x18F       PROTECTED

DATABANK    NAME=gpr0       START=0x20      END=0x7D
SHAREBANK   NAME=dbg0       START=0x7E      END=0x7F        PROTECTED
// This bank is an overlay on the 16F873A so I removed it
// and changed the one bank to DATABANK
// since the memory is still banked (in two banks)
//SHAREBANK  NAME=gpr0      START=0x120     END=0x17D
SHAREBANK   NAME=dbg0       START=0x17E     END=0x17F  PROTECTED

// Same as above
DATABANK    NAME=gpr1       START=0xA0      END=0xF3
SHAREBANK   NAME=dbg1       START=0xF4      END=0xFF        PROTECTED
//SHAREBANK  NAME=gpr1      START=0x1A0     END=0x1F3
SHAREBANK   NAME=dbg1       START=0x1F4     END=0x1FF  PROTECTED

SECTION     NAME=STARTUP    ROM=vectors     // Reset and interrupt vectors
SECTION     NAME=PROG1      ROM=page0       // ROM code space - page0
SECTION     NAME=PROG2      ROM=page1       // ROM code space - page1
SECTION     NAME=IDLOCS     ROM=.idlocs     // ID locations
SECTION     NAME=CONFIG     ROM=.config     // Configuration bits location
SECTION     NAME=DEEPROM    ROM=eedata      // Data EEPROM
```

Notice that the main program imports a delayct variable and a Delay subroutine from Listing 2. This is a simple 16-bit delay loop in a different module. However, since the module is self-contained, it would be just as easy to link it with other programs that need a similar delay. With ordinary methods, you'd have to worry if the new program used any variables the delay module uses and find a place to insert it into the program. With relocatable code, the linker does it all for you.

You can examine what the linker is doing by looking at the MAP and LST files. One thing you'll notice is that the PAGE-SEL and BANKSEL instructions are not always as efficient as they could be. Also, if you lavishly use them before every variable, they may not be needed in every case. To be more efficient, you'd need to make sure certain data sections were grouped into particular banks and take advantage of that fact in your program. Also, judicious use of the UDATA_SHR section can reduce your need for BANKSEL instructions. After all, if data is visible from all banks, there is no reason to select a particular bank before using the variable.

Testing

Even if you don't have any hardware, or you want to apply these techniques to a different PIC chip, you can still try out the examples included. Simply go to the Debugger | Select Tool menu and pick MPLAB Sim as the debugger.

Then you can single step or animate the program, examine variables, and basically check it out without building anything.

What Else?

Once you can make handy reusable modules, you'll probably develop a collection of things you use over and over again. Maybe a software UART, or code to read the A/D converter. That's where MPLIB comes in handy. You can use MPLIB to collect object files into a library. You can add these libraries to your projects and the linker will pull code from the library as needed.

With a little practice, relocatable code becomes second nature. Sure, you probably don't want to use object files when you are trying to squeeze every instruction into the processor (although with careful use of custom linker scripts, you can). However, if you have some room to spare, the development efficiency of using libraries over and over again make it tempting to change your programming style. **NV**

I Love My Heils!

by Dennis Eichenberg

My love affair with Heils has raged on for more than a quarter century! When I stepped into that audio store in 1973 and heard the Heil Air Motion Transformer (AMT) loudspeaker for the first time, I knew that I was hearing music as it was meant to be heard — crystal clear and as smooth as glass! Fortunately, the AMT is still available today, continuing to bring beautiful sound into the world.

What is a Heil?

Dr. Oskar Heil was one of the leading scientists of the 20th century, breaking new ground in a wide variety of physics and engineering fields. Perhaps most notable among his numerous inventions is the Field Effect Transistor (FET), which Dr. Heil patented in 1934. Dr. Heil also did significant research concerning the human ear. He went on to study how small animals were able to produce high sound levels disproportionate to their relatively small size. This research culminated in Dr. Heil's formulation of his basic loudspeaker diaphragm design theory, and the subsequent development of the AMT in 1964. A commercial version of the AMT is shown in Figure 1.

From the early talking machines through the component speakers of the 1960s, sound was produced by a solid material pushing directly against the air creating air movements. The air is only able to move as fast as the pushing solid. The ideal loudspeaker would create sound by using air to move air. By the direct application of the laws of physics, Dr. Heil freed loudspeaker design from the restrictions of mass and inertia, to produce the AMT with its superb transient response and clarity.

A basic diagram of the AMT is shown in Figure 2. The AMT has an extremely lightweight diaphragm, folded into a number of accordion-like pleats to which aluminum strips are bonded. This diaphragm is mounted in an intense magnetic field and the audio signal is applied to the aluminum strips. This causes the pleats to alternately expand and contract in a bellows-like manner in conformance with the audio signal, forcing air out of one side of the pleats, and drawing air into the other. The air movement is five times greater than that of the diaphragm, and the velocity of the air is five times greater than the velocity of the diaphragm. Thus, the device is truly an Air Motion Transformer.

Let's Experiment!

The operating principles of the AMT can be demonstrated by a few simple experiments.

Hold a regular 8-1/2 x 11-inch sheet of paper on one edge with your right hand. Move the sheet back and forth with about a five-inch excursion three times a second. There is very poor movement of the air since you are only transmitting force to the paper from one end, causing the paper to flex and dissipate away the motional energy of your moving arm, instead of transferring it into air motion.

The paper produces substantial noise as it flexes. The motional energy of your arm is largely turned into sound energy, in this case that of the paper's own high Q resonance as the paper flexes, instead of accelerating the air. Moving the paper back and forth less can reduce the noise. Move the sheet back and forth with about a one-inch excursion, thus moving it slower.

Now, the noise generated by the paper is significantly less

Figure 2. An audio signal applied to the conductive strips within the AMT diaphragm open to the front squeezing out air, while the pleats alongside of them, open to the rear, drawing air in.

than with a five-inch excursion. Thus, the lower the diaphragm velocity, the less of its own spurious noise is generated.

Actually, spurious sound generation increases with the square of the diaphragm velocity. Moving the diaphragm half as fast decreases noise generation by one fourth. Moving the diaphragm one fourth as fast decreases noise generation by 1/16th. Thus, a slight decrease in diaphragm velocity provides a significant decrease in noise generation.

Hold the sheet of paper with both hands. Move the paper back and forth. The force of your arm motion is more evenly distributed throughout the paper's entire surface than it is when held by one hand. At both lower and higher velocities, the noise generated by the paper is reduced. Thus, the more evenly the driving force is applied over the entire surface of the diaphragm, the less likely the diaphragm is to flex and generate spurious noise.

Crinkle a sheet of paper, and then crinkle a plastic sandwich bag. Notice how quiet crinkling the plastic bag is compared to crinkling the sheet of paper. The quietest materials have high internal molecular damping. This is because the molecules are of such a shape that they glide over each other rather than crashing into each other. Thus, the higher the internal molecular damping of the diaphragm material, the less of its own noise is generated when it is set in motion.

Another concept incorporated into the design of the AMT, but which cannot be demonstrated through a simple experiment, relates to the energy required to move a diaphragm to a given velocity. The energy required to move any solid is defined by the expression $E = mv^2/2$. Thus, energy increases proportionally to the diaphragm's mass, but geometrically with the square of the diaphragm's velocity. A musical passage requiring a brief but large increase in sound output can require a huge increase in amplifier power, and cause significant distortion. Thus, a lightweight and low velocity diaphragm reduces audio distortion.

The advantages of the AMT are obvious as demonstrated by the performed experiments. The five-to-one transformation ratio of the AMT requires the diaphragm to move at only one-fifth the velocity of the air itself. Since $E = mv^2/2$, the AMT diaphragm absorbs only 1/25th the

motional energy of a conventional loudspeaker. Thus, the AMT has only 1/25th the potential for generating spurious noise as a conventional loudspeaker.

The AMT diaphragm operates by running the audio signal current through a conductor immersed in a magnetic field. The conductor is a very thin strip of aluminum whose width covers almost the entire radiating area of the diaphragm folds. Thus, the resultant motive force generated by current flowing through the aluminum strips is evenly distributed throughout the diaphragm. The more evenly the diaphragm force is applied over the entire surface of the diaphragm, the less likely the diaphragm is to flex and generate spurious noise.

The manner in which the AMT diaphragm is moved permits the use of an extremely soft diagram material. Since the higher internal molecular damping of the AMT diaphragm material generates very little noise when in motion, the AMT is an extremely clean loudspeaker. Since the driving force applied to the diaphragm is distributed evenly across the aluminum conductors laminated on both sides of each fold, there is no tendency for the diaphragm to flex. Consequently, there is no need for a stiff and noisy material to be used to counteract the flexing tendency caused by uneven diaphragm forces.

The unique construction of the AMT provides many additional benefits.

The rise time of a loudspeaker is critical, because the human ear identifies sound in its first milliseconds. The AMT square wave rise time of 15 milliseconds is as fast as the transient resolution of the human ear. The folding of the AMT diaphragm into its transformation alignment shrinks the sound source by a factor of five. This results in horizontal dispersion that exceeds 120 degrees, and provides superb depth of field and excellent horizontal imaging.

The conductors carrying the audio current in the AMT diaphragm are very large and exposed to the air, causing it to be self-cooling. As the AMT is driven harder, it brings a greater amount of air into contact with the conductors. This results in a very rugged speaker capable of developing very high sound levels for a long duration without distortion. An unusual effect of the low distortion of the AMT is that of listener fatigue; the gradual exhaustion as the human ear reacts to subtle distortions in complex tones, is eliminated.

Applications

Loudspeaker systems utilizing the AMT have been very successful. Sources for AMT's, AMT based loudspeaker systems, and AMT replacement parts are indicated in Table 1.

The AMT is normally used as a midrange-tweeter in conjunction with a conventional cone woofer. The ESS AMT Monitor loudspeaker system has an extremely flat frequency

response of 30 to 23,000 Hz, plus or minus 3 dB. It is very efficient, producing an SPL of 93 dB at 1w/m, yet can handle 400 watt musical peaks without distortion.

The AMT has been applied successfully to stereo headphones. The ESS Phase 1 headphones contain a miniature full-range AMT in each earpiece. The frequency response is 20 to 50,000 Hz, with a distortion rating of 0.3% for a 90 dB SPL at 1 kHz.

The ultimate AMT based loudspeaker system ever developed was the Transar. The Transar was an extremely large, full range loudspeaker system developed by Dr. Heil. It consisted of a series of rods and discs in a push-pull arrangement that emitted sound from a dipolar baffle.

Conclusion

Dr. Heil was fond of noting that because nature is complex and truth is manifold, that when truth is achieved, nature is abundantly generous. The characteristics of the Heil Air Motion Transformer: flat frequency response, consistent impedance, fast rise time, high efficiency, wide dispersion, and self cooling, suggests the truth of the design. Thank you Dr. Heil, your legacy lives with me every day of my life. **NV**

DATA PROCESSING USING SCAM

PART 1 — Handling Data Structures

by Gamal Ali Labib, PhD

Many applications require wire speed searches of information. The searches normally involve simultaneous comparison of the desired information against the entire list of prestored entries. Image, voice, computer, and communication systems are possible platforms for such applications.

My proposed SCAM (Super Content-Addressable Memory) may have great potential in supporting DNS operations, URL filtering, network devices address lookup and translation, and virus signature matching.

In the first two articles of this series, I will describe the organization of SCAM and show how it manipulates different data types and queries. The third article will describe some associative logic and mathematical algorithms, and will elaborate on performance issues and calculations relevant to SCAM. The fourth article will be dedicated to the realization of SCAM building blocks using ispLEVER tools. The final article will cover assembling those building blocks into a FPGA device using ABEL hardware description language.

SCAM Overview

Super Content Addressable Memory (SCAM) is intended for integration with general-purpose processors in personal computers and multiprocessor machines. The SCAM module has direct memory access (DMA) to the computer's local memory and has access to the secondary storage via the computer's I/O controller. Such architecture accelerates data movement to/from the SCAM without overhead on the CPU. The SCAM receives data chunks (or data pages) to be processed alongside user queries or operations to be performed. The CPU is then freed to execute other tasks while the SCAM crunches cached data.

As shown in Figure 1, SCAM is composed of three main blocks: the controller (SCAMC), the data repository (SCAMD), and the response block (SCAMR).

The SCAMD incorporates the data cells (denoted by CC: CAM Cell), forming words (which are each eight-bit long) padded with two types of associative cells in each word: the *Structure-delimiter* type (denoted by SC: Structure Cell), used to mark the first word of the data structure header (and, optionally, its trailer); and the *Element-delimiter* type (denoted by EC: Element Cell), used to mark the first word of each data element within the data structure, and to navigate throughout data structures.

The structure-and element-delimiter cells can be manipulated as normal data cells. However, the element-delimiter cell has the additional feature of combining its state outputs in the word control circuitry that incorporates the Tag bit. A memory word may have one element-delimiter cell at the most, but may have more than one structure-delimiter cell. Allocating a structure-delimiter cell to each constituent-structure provides optimum performance for accessing the components of complex-structure. However, a single structure-delimiter cell per word would suffice to minimize the overhead of control gates per associative word, but with increased navigation overhead.

The SCAMC is composed of the Comparand cells (denoted by CE: Comparand Element), the Mask cells (denoted by ME: Mask Element), and the control cells (denoted by MC: Management Cell). The SCAMC controls the cascading of SCAM devices by passing through data lines output from another SCAM device to the SCAMD of its device. It also manages the connection of the SCAM device to the computer chipset to get data and control signals from system buses.

The SCAMR contains the Tag cells that store the results of comparing each SCAMD word against the Mask'ed Comparand word. Tag cells also control forward and backward navigation within the SCAM. Selecting a SCAMD word is accomplished by setting its Tag bit[1] to 1. Two commands affect the setting of this cell: *Compare* and *Set*. The Compare command matches the contents of each of the selected words with the Comparand according to the Mask setting, and, if no match is found, the corresponding Tag bit is reset to 0, or it is left at its current state. The Set command sets all Tag bits in the SCAM to 1. Each Tag cell is linked to its immediate neighbors via its control circuitry that permits propagating the Tag setting to the next/previous data word or element-delimiter word. Four navigational commands allow navigation within the selected structures, those having words with their Tag bits set to 1:

· **link-next-word (LNW)** to select the next-to-current selected word, and de-select the current one.

· **link-next-element**[2] **(LNE)** to select the next-to-current selected element, and de-select the current one.

· **link-previous-word (LPW)** to select the previous-to-current selected word, and de-select the current one.

· **link-previous-element**[2] **(LPE)** to select the previous-to-current selected element and de-select the current one.

With flexibility of navigation within the stored data structures in the forward/backward direction, the predicates in a multiple-element search condition can be evaluated in any order, irrespective of their physical locations within structures.

The SCAM has two modes of operation that determine how the associative words are affected by the launched operations (either navigational or data manipulation):

· The sequential mode causes only the top-most selected word in the SCAM to be affected.

· The parallel mode affects all selected words simultaneously.

Controlling the mode of operation is realized via the *Mode* line input to the SCAM. Intermixing sequential and parallel modes of operation is supported within the same transaction. In the following sections, I shall demonstrate how this design fulfils the requirements stated earlier.

All commands affecting the operation of the SCAM may be issued from assembly or high-level programming language programs that are eventually translated into signals input to the SCAM device(s).

Data Structures and Objects

Here is a quick review of data structures in modern programming languages, including those supporting object-orientation. This review is important for the reader to understand what the SCAM is intended to process.

In the C language, a structure is a collection of variables that are referenced under one name, providing a

convenient means of keeping related information together. A *structure definition* forms a template that may be used to create structure variables. The variables that make up the structure are called *structure elements*.

The example given below is for a person's "address" structure template:

```
struct addr {
    char name[30];
    char street[40];
    char city[20];
    char state[20];
    unsigned long int zip;
};
```

To declare an actual variable with this structure, I would write:

```
struct addr addr_info;
```

This declares a **struct** variable of type **addr** called **addr_info**. The compiler automatically allocates sufficient

FIGURE 1. The SCAM Organization.

Word Contents	Structure-Delimiter	Element-Delimiter	ASCII Data
Structure-ID (SOS)	I	0	7D
Variable-ID	0 0	0 0	FE A0
Name 'J' 'o' 'a' 'n'	0 0 0 0	I 0 0 0	4A 6F 68 6E
Rank 'A	0	I	4I
EOS	I	0	I0

TABLE I. Simple-Structure Format.

memory (shown below) to accommodate all the variables that make up the structure.

```
Name        30 bytes
Street      40 bytes
City        20 bytes
State       20 bytes
ZIP          4 bytes
```

Objects, on the other hand, are the most important feature of an object-oriented language, such as C++. An object is a logical entity that contains both data and code that manipulates the data. To create an object in C++, you must define its general form by using the keyword **class**. A **class** is similar syntactically to a structure. The following **class** defines a type called **addr**, which is used to create and manipulate a person's addresses within a user program.

```
class addr {
    char name[30];
    char street[40];
    char city[20];
    char state[20];
    unsigned long int zip;
public:
    void init();
    void set_name(char *n);
    void set_street(char *n);
    void set_city(char *n);
    void set_state(char *n);
    void set_zip(long i);
};
```

To create an **object** (or an "**instance**") called **addr_info** of class **addr**; use the following line of code in the program:

```
addr addr_info;
```

Data Manipulation Requirements

Querying composite objects or nested structures (i.e.,

those referencing or containing instances of other classes/data types) is accomplished through the following sequence of operations:

Restricting (or selecting) composite and constituent data as specified by the query search conditions.

If constituent data is not clustered (e.g., adjacent in memory), data may need to be joined where a nested-structure is merged with its referenced-structures and the latter are merged with their basic data elements.

Associative techniques, utilizing SCAM, can be advantageous in improving the execution time of restriction and join operations. Accordingly, take into account the following basic tasks in the SCAM design:

• Identifying data structure boundaries (e.g., structure header and trailer if used by the user application). This aims at reducing the control information stored in the data-page header and the overhead involved in processing such information and locating constituent structures.

• Locating data elements within structures. This aims at reducing the control information stored in the structure header and the overhead involved in locating specified elements.

• Handling multiword data elements. This involves the manipulation of both numeric and alphanumeric data (such as text strings of variable length).

• Moving activity between neighboring data words. This provides for accessing specific words within data elements. The benefit of this is obvious when processing particular control information in the structure header, or when seeking partial element matches.

• Moving activity between neighboring data elements. This provides for processing more than one predicate involved in the same query, without the need to store intermediate results.

• Parallel-by-word operation. For this operation, the ultimate performance can only be achieved by SCAM organization.

The current content-addressable memory (CAM)

	Word Contents	Structure-Delimiter	Element-Delimiter	ASCII Data
Employee Composite Structure	Structure-ID (SOS)	I	0	7D
	Variable-ID	0 0	0 0	FE A0
	Name 'J' 'o' 'h' 'n'	0 0 0 0	I 0 0 0	4A 6F 68 6E
	Rank 'A	0	I	41
	Spouse-ID	0 0	I 0	FE B0
	Child-ID	0 0	I 0	FE C0
	EOS	I	0	10
Spouse Constituent Structure	Structure-ID (SOS)	I	0	7E
	Variable-ID	0 0	0 0	FE B0
	Name 'J' 'o' 'a' 'n'	0 0 0 0	I 0 0 0	4A 6F 68 6E
	Age	0	I	1F
	EOS	I	0	10
Child Constituent Structure	Structure-ID (SOS)	I	0	7F
	Variable-ID	0 0	0 0	FE C0
	Name 'D' 'a' 'v' 'e'	0 0 0 0	I 0 0 0	4A 6F 68 6E
	Age	0	I	0F
	EOS	I	0	10

TABLE 2. Composite-Structure Format in Breadth-First.

architectures lack the hardware support for all those tasks in the same design, and in attempting to achieve some of them, other drawbacks emerge. For example, some designs impose a restriction on data by reserving a specific bit pattern for data element headers, others limit

RESOURCES

1. The SCAM's Tag bit corresponds to that in conventional **AM** but with additional control circuitry.

2. This control command is not effective if the current selected word is an element-delimiter.

3. May also represent Object-ID if you are working in an object-oriented environment.

4. Value may vary depending on user choice.

	Word Contents	Structure-Delimiter	Element-Delimiter	ASCII Data
Employee Composite Structure	Structure-ID (SOS)	1	0	7D
	Variable-ID	0 0	0 0	FE A0
	Name 'J' 'o' 'h' 'n'	0 0 0 0	1 0 0 0	4A 6F 68 6E
	Rank 'A	0	1	41
Spouse Constituent Structure	Structure-ID (SOS)	1	0	7E
	Variable-ID	0 0	0 0	FE B0
	Name 'J' 'o' 'a' 'n' Age	0 0 0 0 0	1 0 0 0 1	4A 6F 68 6E 1F
	EOS	1	0	10
Child Constituent Structure	Structure-ID (SOS)	1	0	7F
	Variable-ID	0 0	0 0	FE C0
	Name 'D' 'a' 'v' 'e' Age	0 0 0 0 0	1 0 0 0 1	4A 6F 68 6E 0F
	EOS	1	0	10
	EOS	1	0	10

TABLE 3. Composite-Structure Format in Depth-First.

the movement of activity to one direction only, and many do not implement true parallel-by-word operation.

Data Representation

The information stored in data structures mainly depends on the user application needs. However, the delimiter words of the structure header/trailer and subsequent data elements should be marked using the structure-delimiter and the element-delimiter bits. Accordingly, there would be no need for embedding control information in the data structure to locate such boundary words.

In our case, I choose to record the Structure-ID in the structure header, which I call the **S**tart-**O**f-**S**tructure (SOS) word. The SOS would have its structure-delimiter bit set and its data word loaded with ASCII '7D' corresponding to the Structure-ID value of 125. I follow the SOS word with one or more words for the Variable-ID[3], which has its control bits reset.

For subsequent structure elements, the first word of an element should have its element-delimiter bit set. Finally, I assign an **E**nd-**O**f-**S**tructure (EOS) word to the structure with its structure-delimiter bit set and its data word loaded with ASCII '10'[4]. Table 1 illustrates the word settings of a simple-structure.

The representation of clustered composite-structures is similar to that of simple-structures, where each constituent-structure is formatted as if it were a simple-structure on its own, but in a contiguous space to the composite-structure. If constituent structures are stored in breadth-first (see Table 2), then their composite-structure would have a Variable-ID element representing each of those structures.

If composite-structures are stored in depth-first (see Table 3), then constituent Variable-IDs are removed from the composite-structure and the EOS word of the latter follows the last constituent-structure. The representation of un-clustered composite-structures would be similar to that of breadth-first, while constituent-structures may be located anywhere in their database and not restricted to the vicinity of their composite-structures.

To complement this description of the SCAM operation, next month's article will demonstrate how data manipulation and navigation commands are applied while resolving queries targeting simple and complex data objects. **NV**

Advanced Techniques for Design Engineers

The Design Cycle

The SpectroTech SpectroBUS18

If you're a PICBASIC programmer, this month's spin of Design Cycle is for you. I came across some rather unique prototyping hardware that happens to have microEngineering Labs' PICBASIC at its core. The PICBASIC-based system is called SpectroBUS18. There are plenty of goodies in the SpectroTech hardware package, so let's get started.

The SpectroTech SpectroBUS18

The SpectroBUS18 prototyping system is a collection of dumb and/or smart modules that interact via a system of busses that reside on the main motherboard (Photo 1). The SpectroBUS18 bus architecture is such that any module adhering to the SpectroBUS18 bus can be plugged into any of the SpectroBUS18 motherboard's six positions (A-F).

This scheme allows a small PIC microcontroller with a complement of only eight-bit I/O ports like the PIC16F88, which supports only PORTA and PORTB, to logically look like a high-pin-count PIC device with multiple I/O ports.

In addition to supporting multiple devices via a data bus, the SpectroBUS18 also includes an auxiliary bus that extends smart module reset signals, the MCLR signal, module oscillator signals, debug signals, and module power.

The idea behind the SpectroBUS18 is proof-of-concept prototyping. However, the SpectroBUS18 is compact and powerful enough to also be used in a production environment. The SpectroBUS18 is designed to allow you — the system designer — to apply a building block approach to the design of the hardware.

Each SpectroTech smart module is supported by an associated building block of skeleton source code in PICBASIC format. The absence of schematic diagrams for any of the SpectroBUS18 system components was at first a great shock, as I always look at the schematic diagrams to help make sense of the hardware and supporting firmware.

However, after carefully studying the SpectroBUS18 motherboard layout and reading the SpectroTech smart module documentation, I found that I really didn't need a schematic diagram of each SpectroBUS18 module to apply the smart modules and SpectroBUS18 motherboard to an application.

If you're interested in multi-microcontroller applications, the SpectroBUS18 system is a good example of how to take advantage of multiple intercommunicating micro-controllers. Each of the SpectroBUS18 smart modules is controlled by its own PIC microcontroller, which also performs the communications tasks between the smart module I/O and the SpectroBUS18 motherboard's resident PIC16F88.

My SpectroBUS18 prototyping system came complete with three smart modules, an ICSP programming adapter, a couple of SpectroBUS18 breadboards, a ZIF adapter, and numerous connector interface modules.

The SpectroBUS18 system assumes that you will be programming with microEngineering Labs' PICBASIC or PICBASIC Pro. Although the SpectroBUS18 documentation states that the SpectroBUS18 system supports a number of PIC programmers, the SpectroTech hardware and documentation lean towards PIC programmers that can be purchased from microEngineering Labs. The EDTP Electronics library contains a copy of PICBASIC Pro but, unfortunately, there were no microEngineering Labs PIC programmers on the shelves. That's where the fun begins.

Enabling the Microchip ICD2

PICBASIC Pro interleaves into the MPLAB IDE seamlessly, and having MPLAB IDE in the PICBASIC Pro mix

Photo 1. The SpectroBUS18 motherboard is actually a complete single board computer complete with a PIC16F88, a supporting crystal oscillator circuit, ICSP programming port, and a regulated +5VDC power supply.

makes using the MPLAB ICD2 programmer/debugger with PICBASIC Pro applications a logical choice. I've developed with the MPLAB ICD2/PicBasic Pro/MPLAB IDE combination in the past, and I fabricated a custom cable that allows me to interface the MPLAB ICD2's six-pin RJ-12 to the microEngineering Labs protoboards' AVR-like 10-pin programming interface. The details of my custom MPLAB ICD2-to-microEngineering Labs programming/debugging cable are shown in Figure 1.

Figure 1. Since the MPLAB ICD2 connector and the target programming/debugging receptacle are both female, the MPLAB ICD2 programming/debugging interface cable is assembled in a crossover configuration.

Note that both the MPLAB ICD2 and MPLAB ICD2-compatible target programming/debugging sockets are terminated with a six-pin RJ-12 female connector. That means we have to construct the MPLAB ICD2 programming/debugging cable with a crossover configuration. This is easily done by simply crimping on the RJ-12 male ends with the gold pins facing up at each end of the cable. The resulting crossed-over connections at the opposite end of the MPLAB ICD2 female RJ-12 connector are then fitted to match the microEngineering Labs 10-pin configuration. With that said, the MPLAB ICD2 with its custom microEngineering Labs-style programming/debugging cable should interface without a hitch to the SpectroBUS18's ICSP adapter.

That would have been too easy. After I connected the MPLAB ICD2 to the SpectroBUS18 motherboard's ICSP adapter, I fired up the MPLAB IDE and encountered a MPLAB ICD2 self test failure. Examining the MPLAB ICD2's status revealed a low target Vdd voltage. Some probing with a voltmeter led to the absence of +5.0 VDC on pin 1 of J1 on the SpectroBUS18 ICSP adapter (the MPLAB ICD2 side).

A jumper wire between a +5.0 VDC point on the SpectroBUS18 ICSP adapter to pin 1 of the ICSP adapter's J1 header solved the MPLAB ICD2's low Vdd problem. Celebrating over what I thought would be a quick and nasty fix, I once again fired up the MPLAB IDE. This time I did not get any MPLAB ICD2 failures. So, I compiled some code and attempted to program the resultant hex file into the PIC16F88 on the SpectroBUS18 motherboard. The program operation was a success. However, I could not get the SpectroBUS18 system to run with the MPLAB ICD2 attached to the SpectroBUS18 ICSP adapter.

That's not good, as the SpectroBUS18 documentation says that the microEngineering Labs PIC programmers don't have to be removed from the circuit to run the SpectroBUS18 system after programming. And, when using the microEngineering Labs PIC programmers, the SpectroBUS18 ICSP adapter is smart enough to know when to send a system reset signal to the SpectroBUS18 motherboard and all of the smart modules in the system when the programming process is completed. At this point, I figured it was time to get to know the SpectroBUS18 ICSP adapter intimately.

Schematic 1. This is pretty darned clever. The PN2222 is, in effect, a high-voltage TTL buffer that feeds a pulse generator in the guise of a PIC12F629. I modified the original circuit to accommodate my MPLAB ICD2.

Photo 2. I used a Moto Tool to cut the PGC and PGD traces between J1 and J2. I then added a couple of 1K SMT resistors across my cuts in the traces. A simple jumper wire between pin 1 of the PIC12F629 and pin 1 of J1 cured the MPLAB ICD2's low Vdd problem.

The SpectroBUS18 ICSP Adapter

I've provided a schematic view of the SpectroBUS18 ICSP adapter in Schematic 1. As you can see in the SpectroBUS18 ICSP schematic, I made a trio of modifications. If you look closely at Photo 2, you can see that I added 1K SMT resistors in series with the PGD and PGC PIC programming lines between the MPLAB ICD2 ICSP connector and the SpectroBUS18 ICSP connector.

The sole purpose of the SpectroBUS18 ICSP adapter is to generate a reset pulse on the SpectroBUS18's SPDR line and pro-

Photo 3. VDD and GND are on the socket's odd pins, while all of the other bus signals are positioned on the socket's even pin numbers. On the data bus side, PORTA consumes the odd pins and PORTB occupies the even pins. Despite the absence of schematics, the SpectroBUS18 documentation clearly lays out the pinouts of all of the SpectroBUS18 system components and connectors.

vide an interface between the larger standard .1 in-center 10-pin header used by the microEngineering Labs PIC programmers and the smaller 2 mm ICSP interface on the SpectroBUS18.

The SPDR line is an active-low bus signal that is connected to each smart module's MCLR line via the module sockets (A-F) on the SpectroBUS18 motherboard. Photo 3 provides the visual details of a pair of the SpectroBUS18 bus sockets. The motherboard's PIC16F88 MCLR line is not reset by the SPDR signal. Instead, the PIC programmer attached to the SpectroBUS18's ICSP connector provides a reset signal for the PIC16F88 in a normal manner when programming is completed.

Each of the SpectroBUS18's smart modules includes a status LED that illuminates and then extinguishes when the module has been addressed and initialized successfully. I noted after programming the SpectroBUS18 with the MPLAB ICD2 that I could toggle the PIC16F88's MCLR line via the MPLAB IDE's Hold/Release Reset commands and each of the smart module's status LEDs would illuminate and extinguish in the order of their addressing.

However, the program would not run. I also noticed that, after the initial release from reset operation, toggling the Hold Reset and Release Reset within MPLAB IDE would not cause a toggle of the smart module status LEDs and the program still would not run. Sometimes I got blinking status LEDs, which was bad, and sometimes I didn't. Nothing was failing in a consistent manner. I was baffled. So, I decided to stick a logic probe on

the SPDR line to see what it was doing and when it was doing it.

I found that the SPDR line is always held at a logically high state when the applied Vpp (via the MPLAB ICD2 at J1 pin 2) is at 0.0 volts. Applying +5.0 VDC to J1's Vpp pin also yielded a high SPDR line. When +13.5 VDC was presented at J1's Vpp pin, a low-going pulse is emitted from the SPDR line when the Vpp voltage falls from +13.5 VDC to below +5.0 VDC. I concluded that, to initiate a reset of the SpectroBUS18's smart modules via the SPDR line, the Vpp voltage must rise to +13.5 VDC and then fall back to 0.0 VDC to initiate a low-going pulse from the SpectroBUS18 ICSP adapter's SPDR pin.

That explained why the MPLAB IDE Hold/Release reset commands didn't work. The MPLAB IDE Hold/Reset command only toggles the level of the PIC16F88's MCLR pin between 0.0 VDC and +5.0 VDC. Since no +13.5 VDC is generated by the MPLAB ICD2 when the MPLAB IDE Hold/Release commands are used, the low-going SPDR pulse is never generated by the PIC12C629 and the smart modules don't get reset via the SPDR line.

Here's how the low-going pulse is produced by the SpectroBUS18 ICSP adapter: The MPLAB ICD2 actually produces +13.5 VDC during programming. When the voltage at J1's Vpp pin is below +13.5VDC, a TTL low (less than +0.8 VDC) exists at the emitter of the PN2222 transistor. A TTL high (greater than +2.0 VDC) at the emitter of the PN2222 transistor only occurs when the voltage at J1 pin 2 reaches +13.5 VDC. The voltage at the emitter of the PN2222 transistor, which is a TTL compatible voltage level, is being read by the GP3 pin of the PIC12F629. The three resistors supporting the PN2222 are selected to swing the PN2222's emitter voltage between near 0.0 VDC to just over +2.0 VDC.

The PIC12F629 specifications state that a voltage of +2.0 VDC is the minimum required on any input pin to be interpreted as a TTL high level.

NUTS & VOLTS

Since we're not privy to the PIC12F629 firmware, we can only assume that maybe the weak pullups are enabled on the GP3 pin to insure the detection of a TTL high at the GP3 pin. What we do know about the firmware inside the PIC12F629 is that it is looking for a TTL-low-to-+13.5 VDC-to-TTL-low transition before producing a low-going pulse at the PIC12F629's GP4 pin, which is feeding the SpectroBUS18's SPDR line via a 100-ohm resistor.

Not being privy to the microEngineering Labs' PIC programmer schematics, I must assume that, to remain attached to the SpectroBUS18' ICSP connector after programming, microEngineering Labs' PIC programmers include integral isolation circuitry for the PGC and PGD programming lines, which are PORTB pins 6 and 7, respectively.

Since the PGC and PGD programming lines are actually part of the PIC16F88's PORTB I/O port, I added a couple of 1K resistors between MPLAB ICD2 and the PGC and PGD pins of the PIC16F88 that resides on the SpectroBUS18 motherboard. Adding the isolation resistors to the SpectroBUS18 ICSP adapter hardware allows the MPLAB ICD2 to remain connected to the SpectroBUS18's ICSP connector by eliminating any interference that the MPLAB ICD2 circuitry may cause to the PIC16F88' PGC and PGD lines, which are used as I/O port and bus lines when the SpectroBUS18's PIC16F88 is not in program mode.

With the MPLAB ICD2 problems solved, I was able to compile, program, and run a PICBASIC Pro program with an MPLAB ICD2, just as if a microEngineering Labs' PIC programmer was attached. Please don't misconstrue my intentions here. If you have a microEngineering Labs programmer, it's a good product. Use it. I just didn't have one handy and had to come up with a workaround for my MPLAB ICD2.

That pretty much covers the SpectroBUS18 motherboard and the SpectroBUS18 ICSP adapter. Let's move on and look at the SpectroTech smart modules.

The SpectroBUS18 Smart Modules

Bear in mind that the engineers at SpectroTech know that you and I can write code and build hardware that processes keystrokes, drives an LCD, or sends an RS-232 message. The thinking behind the SpectroBUS18 smart module technology is that, if you have a concept in mind, you can eliminate the time it takes to build up the basic building blocks of hardware before laying down the code for your design. Also, since the building blocks are actually removable modules, you can recreate any past design by simply plugging in the correct modules and reprogramming the SpectroBUS18's PIC16F88.

SpectroBUS18 smart modules communicate with the host PIC16F88 serially via the bus using PORTB pins 6 and 7. When the bus is run on the PORTB side of the SpectroBUS18 motherboard, each smart module is addressed with the remaining bits of PORTB. For

Photo 4. The tri-state address DIP switches are prominent on all three of these smart modules. I'm sure you can imagine what's running inside each of the smart module PICs. The smart modules you see in this shot are building blocks that you can connect and activate with a minimal amount of coding.

instance, smart module 0 is addressed via PORTB bit 0, smart module 1 is addressed via

Photo 4 is a shot of the three smart modules that came with my SpectroBUS18 system. All smart

Figure 2. The tri-state DIP switch is simple to understand. The eight switch I/O points can be switched between a pair of busses ("+" and "-") or set to a not-connected position ("o"). If you look at the tri-state DIP switch settings in Photo 5, you'll see that the address true setting is on the "-" switch bus and that one of the data communications signals (DIP switch position 2) is transferred on the "+" switch bus. The data communications signal at DIP switch position 1 is directly connected to the smart module's PIC. Therefore, DIP switch position 1 is always placed in an "o", or not-connected state.

modules are physically addressed using an eight-position tri-state DIPswitch, which can be seen photographically on the smart modules in Photo 4 and schematically in Figure 2. Positions 1 and 2 of the tri-state DIP switch, which relate directly to PORTB bits 7 and 6 on the bus, route serial data and communications signals between the smart module and the host microcontroller on the SpectroBUS18 motherboard.

Positions 3 through 8 of the tri-state DIP switch are switched onto the bus depending upon the smart module's physical address. The tri-state DIP switch can be physically moved to the smart module's PORTA position. However, there is one less smart module

physical address on the PORTA bus. And, once a smart module occupies a bus, the I/O port that is involved with that particular bus cannot be used for standard I/O operations by the PIC16F88. If your application so requires, it's a good idea to leave one of the busses open to allow the PIC16F88 to do standard pin I/O.

It's not important to understand a smart module's hardware, as we don't have any information on the firmware running inside the smart module's PIC. However, you can get a good feel as to what a SpectroBUS18 smart module does by just looking at it. Photo 5 is a nose-to-module view of the SpectroCOM Intelligent RS-232 Serial Communications Controller module.

The SpectroCom smart module's PIC16F819 is expected, since all SpectroTech smart modules will house a PIC of some type. And, there's nothing new about the Sipex SP232A and its collection of charge pump capacitors. The only difference in this smart module and something similar that you would design and build is that you didn't have to build it and you didn't have to write code for it. That's the main idea behind the

SpectroTech SpectroBUS18 system.

If you take another look at Photo 4, you'll see that a PIC on each smart module is actually handling the device I/O. This is exactly the same multi-processing method that is employed by mainframe computers and their intelligent peripherals. The mainframe (in our case, the SpectroBUS18 motherboard) sends data and commands to the peripherals (our smart modules) and moves on to something else. The mainframe peripherals and the SpectroBUS18 smart modules are computing devices that are capable of carrying out the instructions from their commanding host without any intervention from the host after the command has been issued.

In the case of the SpectroCom module, the PIC16F819 buffers data to and from the RS-232 serial port, allowing the PIC16F88 on the SpectroBUS18 motherboard to communicate serially via RS-232 without missing incoming characters and while servicing other smart modules that may be attached to the SpectroBUS18 motherboard. Now that you've got a good handle on the SpectroBUS18 hardware, let's put some characters on an LCD using the SpectroLCD module.

Coding the SpectroLCD Module

A comprehensive and well-commented set of skeleton code is provided with the SpectroTech SpectroBUS18 system. To save some space and to make the LCD code a bit easier to read, I've eliminated some of the boilerplate in Listing 1 (See Listing 1 on the *Nuts & Volts* website; **www.nutsvolts.com**). I've also taken the liberty to include the skeleton "INCLUDE" module code for clarity.

For instance, at the beginning of the LCD source code, I've commented out the original "INCLUDE" statement and replaced it with the actual code contained within the included file. In this case, the *include "INC_H_16F88_HS20.BAS"* statement is expanded to reveal that this include code is no more than code that sets up the PIC16F88

fuses. The code within *include* "INC_H_16F88_HS20.BAS" also contains a couple of PICBASIC compiler directives.

Following the code downward shows us nothing that, as PICBASIC programmers, we haven't seen before. However, note the variable definitions for PORTB. Recall that PORTB bits 6 and 7 are used by the smart modules to communicate via the SpectroBUS18 motherboard bus to the host PIC16F88. With these PORTB definitions in mind and examining the source code in Listing 2 (See Listing 2 on the *Nuts & Volts* website), you can get a good sense of how that communications process works.

As you can see in Listing 2, the *Read_Device* and *Write_Device* subroutines use the PORTB bits 6 and 7 as a bidirectional handshaking serial port between the smart module PIC and the SpectroBUS18 host PIC16F88. Our LCD code will only use the *Write_Device* subroutine, as we will only be sending data and commands to the SpectroLCD smart module.

The next *INCLUDE* statement, *INCLUDE* "INC_H_BusSpeed.BAS," is actually a collection of SMODE definitions representing differing baud rates and bus speeds. You simply uncomment the SMODE definition you want to use. Here, I've chosen to hang with the default 9600 bps as a bus speed. The bus speed is the speed at which the PIC16F88 and the smart modules communicate using the PORTB bits 6 and 7 we discussed earlier.

I have jumpered my SpectroLCD module for address 0. The logical configuration is done using the Device Configuration Control Word (DCCW) and the Device Configuration Parameters A, B, and C. These words are transmitted to the smart module during the status LED illumination event I talked about in the MPLAB ICD2 segment of the column and are

stored in the smart module's on-board PIC. A detailed explanation of each bit of the configuration words is included in the SpectroBUS18 documentation. As you can see here, I've settled for the factory defaults for this little application. The configuration words are transmitted by code within the "INC_H_Device_Config.BAS" INCLUDE file. Nothing tricky there.

Writing to the SpectroLCD smart module is as simple as loading up a buffer with our ASCII characters and LCD commands, designating the length of the buffer contents and calling the *Write_Device* subroutine. In fact, that's how all of the smart modules work. Note the use of the buffer in both the *Read_Device* and *Write_Device* source code (Listing 2).

The SpectroLCD system firmware also includes some handy routines and character sets that can be called into use. Notice that a custom Bargraph character included with the SpectroBUS18 software package was loaded right after the *Nuts & Volts* message was sent to the SpectroLCD smart module. The remainder of the

Photo 6. Much can be learned about the SpectroBUS18 and the multi-processing techniques it uses by probing the bus using the screw-down module you see here mounted along with the SpectroLCD. The pot is acting as a voltage divider and is feeding the analog-to-digital converter of the SpectroBUS18's PIC16F88.

code is simply an analog-to-digital converter read operation that displays a bargraph relating to the voltage at the PIC16F88's analog-to-digital converter input. Nothing that you can't handle.

SpectroKickStart

I'll leave you with the results of the LCD code we've just discussed in Photo 6. The SpectroBUS18 system is a cleverly designed development tool. And, the SpectroBUS18 system is a good "quick and nasty" way of accelerating your Design Cycle. **NV**

Sources

SpectroBUS18: **www.spectrotech.net**

PicBasic and PicBasic Pro:
www.melabs.com

CLASSIFIEDS

News Bytes

Impact of Katrina on Shuttle Program

As NASA continues to assess the impact of Hurricane Katrina on the future of the shuttle program, at least one official is warning it could take up to a year before the next flight takes off.

The space agency grounded future shuttle flights after a fuel tank insulation problem was found during Discovery's last mission. The pre-Katrina hope was for a new shuttle flight in March 2006, but after NASA's Michoud facility in New Orleans was hit during the hurricane, analysts expected that mission to slip into May. That now may turn out to be overly optimistic — and not just because of the hurricane damage.

Backpack Generates Its Own Electricity

In the days immediately following Hurricane Katrina, radio and phone communication suffered, in part, when rechargeable batteries died and could not be recharged due to widespread power outages. A new backpack design may offer a way for first responders and disaster relief workers to generate their own electricity for communications devices, night vision goggles, water purifiers, or other crucial, portable electronics.

All the person wearing the backpack has to do is walk — the backpack does the rest. The backpack captures energy from the up-and-down movements of its heavy contents and converts this energy to electricity.

The new research is published in the Sept. 19 issue of the journal *Science,* published by AAAS, the nonprofit science society.

Online Photography Lesson Series

Manfrotto, the world's leading manufacturer of support products for pro video and photography, announces the launch of a series of practical photography lessons available to the public at **www.manfrotto.com** The lessons were developed in collaboration with Web Photo School, the California-based digital photo education company.

Developed to help Manfrotto customers get the most out of their equipment in various "real-life" photographic situations, the lessons provide valuable information on how to tackle different subjects from still life to macro, from nature photography to portraits and more.

Each lesson has been created utilizing a range of Manfrotto tripods and heads. Each lesson provides detailed information on every step of the photographic process, from setting up the tripod through framing and lighting technique news turning a good photograph into a great one.

The Business of Electronics Through Practical Design and Lessons Learned

In The Trenches
Hiring and Firing

As engineering departments grow, engineers are often promoted to management positions, or at least assume management responsibilities. When those responsibilities include hiring engineers and technicians, new managers can be challenged by the difficult assignment. Even more challenging are the times when workers must be eliminated from the department because of financial concerns, or poor performance or behavior. This month, we'll look at the factors involved in hiring and firing.

Why Do You Need Help?

The first step in hiring is to define why additional personnel are needed. Is it because the department is growing? Is it because there are new markets to address? It is to replace someone who has left? Is it because the project is going badly and help is needed to meet a deadline? The circumstances of the hire often determine the qualities needed in an applicant.

Hiring for Growth or Replacement

Adding personnel to a growing department is the easiest and most straightforward hiring task. Because the department has been in existence for some time and the people are used to working together, the skills required are well defined. The new hire should have a personality compatible with others in the department, along with the requisite skills. There is plenty of time available to show the new person how the department


```
SEE HERE;(THIS IS A COMMENT)
          ; Teach your kid that American design needs to be much, much better than
          ; "Good enough for government work." Perhaps the result will be industries
          ; and engineers that can design a Humm_, a LEVEE, or even a Robot.
          ; BNF simple example
          ; Begin
  IV-KITS  ::=      BREADBOARDS
           and      PARTS
           and      CHIPS
           and      CODE
           and      CONSTRAINTS
           where    BREADBOARDS ::= BB2 : BB3 : BB4 : BB5

           --       LOW COST AND SHORT SCHEDULE
           ; end
  ORDER:   "Using RNF for Design" .................. (120 pages softcover) $5.00 postpaid
  USE RNF: as a design language for systems, hardware, software, firmware, packaging,
           debugging, testing, documentation, analysis, synthesis, verification, and
           validation. It is both top_down and bottom_up definitional methodology.
  IV, POB 245, WALDWICK, NJ, 07463-0245
```

works and to allow him or her to learn how things are done.

If the company wants to expand into new areas, it's more difficult to determine the ideal candidate. Because the company has little or no experience with the new market, it's hard to judge the technical skills needed to fill the niche. Additionally, the performance of the person who fills this key position will in large part determine whether or not the expansion will succeed. In this instance, you need someone with initiative, experience, and self-reliance.

When you are replacing someone who has left the company, the job requirements are very specific. If the person left unexpectedly or was in a key position, the need to fill the position can be very urgent. In this case, you want a seasoned applicant who can get up to speed and be productive very quickly.

Hiring to Meet Deadlines

All too often, a company needs to hire people in order to meet contract deadlines that are slipping. This is by far the worst situation under which to recruit new employees. The problem often dates back to the contract proposal, when management made unrealistic assumptions regarding the amount of work (or time) needed for the project or failed to realize that the workload would exceed the available staff. Admittedly, sometimes this is necessary in order to win the contract, but the inevitable results is that overworked and stressed-out employees cannot perform at their best.

Hiring people in order to meet a deadline will actually result in a net decrease in productivity for several weeks. Integrating new people into a department requires the existing (overworked) people to teach the new

hires how the department works, how the job is done, what paperwork is needed, and how to fill it out, etc. This takes valuable time from the already overworked staff, and results in greater staff expenses without a recognizable increase in productivity.

Most of the time, "throwing bodies" at a project simply isn't a desirable solution, particularly when project tasks are already properly staffed. For example, only one person can lay out a printed circuit board. You can't assign the top and bottom layers to two people and expect anything good to happen.

But, if you have to hire new people, determine the specific tasks to which they will be assigned and look for candidates who can learn very quickly and have a very specific skill set. Realize that they will most likely require an above average salary, too. Lastly, it's important to consider what happens to these people after the contract is completed. Can the company afford to maintain the higher staffing level? If not, fewer people will apply the next time you have a crisis, because they will know it's probably just a temporary job.

Job Description

When advertising a job opening, it's important to define the position and outline the responsibilities and skills required. "A BSEE with 2-4 years experience in digital design and good communication skills" is not an adequate description. You'll receive many resumes, but you aren't providing the filter you need to attract the people you want and discourage the applicants you don't want.

The job description should address the character of the position. "Growing company needs an engineer with several years experience in memory interface design. Experience in meeting and discussing projects with clients is necessary." Although this advertisement will yield fewer applicants, those who do apply will better fit your needs.

At the other extreme are the advertisements that are so specific as

to eliminate many qualified applicants: "Must have experience in TeknoCAD Version 3.14159." This is just a silly requirement. Virtually all CAD (Computer Aided Design) systems are the same. If you know one system, you can easily transfer your knowledge to another system. Just as when driving a car, the manufacturer of the car isn't important; what's important is whether or not you can drive well. Because virtually every engineer uses CAD to some degree, it's better to specify the type of CAD required, such as "3-D Modeling" or "Finite Element Analysis."

Filtering

All resumes get filtered in some manner. Either Human Resources filters out those that don't meet the job description, software eliminates emailed resumes that don't contain the proper keywords, or you plow through a stack of resumes and separate them into various piles. Filtering is necessary, but unless the proper filters are in place, you can miss qualified candidates.

As we saw above, your first filter should be your job description and advertisement. If the skills needed are defined too broadly, you will have too many candidates. If they are defined too narrowly, you will filter out desirable applicants. A good job description is a self-filtering mechanism for applicants.

Finding the Best Candidate

Truthfully, there are many people who can perform any given job satisfactorily. Quite often, it's appropriate to place an average person in a general position. Choosing an average person to perform an average job is a fairly easy task.

But, if you want the "best" person for the job, you must expend much more effort. Naturally, the first task is to define "best." Clearly, different people will have different ideas about what constitutes "best," so it's important to specify those attributes that

your "best" candidate must possess.

Suppose you want someone to "build and direct a new research and development group to create innovative products" for the company. This is a critically important position, and you are searching for one special person. Unfortunately, in situations such as these, human nature often makes the job description and filtering

process mutually exclusive. That is, the "best" candidate (the one who best matches the job description) cannot be chosen (is filtered out).

Human nature comes into play because the position is extremely important to the company and the company doesn't want to choose a risky applicant. Instead, they want someone who is conservative with a

good, solid history. However, you will have a hard time finding someone who is both innovative and conventional.

Innovation demands looking at things in new ways and doing things differently. The people who are successful at innovation cannot and will not be conservative. It is their nature to contest the status quo. Typically, these candidates have an unconventional job history. This can cause them to be disqualified — even though they probably have the greatest chance of creating exactly what the company needs.

This example illustrates why it is important to properly measure the needs of the company against the strengths of the applicant, as well as to be consistent with the job description and the evaluation/filtering process. It's also important to understand that a component of the job requirements may include the applicant's personality and behavior. If this awareness isn't present, the best applicant may never be selected.

Person or Paper

Many people fail to understand that they are hiring a person, not a piece of paper. It doesn't matter if the person has a Ph.D. and a perfect résumé. If he or she can't do the job or can't work with others, that candidate is a poor choice. Employment ads often say, "Must have a BSEE." Why?

There are two basic reasons for educational requirements. The first is credibility. Someone with a BSEE degree has shown that he or she has a specific amount of engineering training, so the applicant can be believed when he or she says he or she can do the job. The second reason is that it makes the screening process easier.

If the applicant doesn't list a BSEE on their resume, it can be thrown out. A third reason that sometimes comes into play is that some employers believe that an engineer without a degree might create problems in a department where everyone else has a BSEE. Most engineers, however, judge others by their performance, not their credentials.

If someone has demonstrated the capacity to perform at a high level without a degree, it shows that the person is highly motivated and very capable. Those people are generally good workers. From a pragmatic viewpoint, they will also often work for a lower salary. However, finding such employees means that you have to look past the paper and at the person. It does take effort, but the results can be worth it. Think of it as an Easter egg hunt.

The Cover Letter

Many times, Human Resources (or the employment agency) will remove an applicant's cover letter and only forward the actual résumé. The cover letter is significant, as it tells you about the person. From their cover letter, you can see how well the applicant communicates, what he or she thinks is meaningful, and how he or she approaches a problem. A conventional letter indicates a conventional applicant, while a creative cover letter signifies a creative applicant. The resume is the "paper," while the cover letter is the "person."

The Interview

Unfortunately, the interview process is often given short shrift. Typically, there are one or two meetings that last a few hours each. From this, a decision is made. If successful, the applicant will spend more time at the office with his co-workers than he will with his family (excluding sleep). This is akin to proposing marriage after two dates.

The real purpose of the interview(s) is to determine the compatibility of the applicant with the company. Again, this is an examination of the "person," not the "paper." Assuming

that the résumé is truthful, it should be clear what his or her technical strengths and weaknesses are. Of course, it's useful to determine if the person has "fudged" or "padded" his or her résumé. So, some technical questions are very appropriate. But the interview is not really a test of capability. Rather, it is a test of character. Does the candidate have those personal qualities that are important to the company?

This means that the interviewer should be an able judge of people. This is not always the case for engineers, who (as a group) are often more interested in technical prowess than personality. So, when you are interviewing someone, think about what it would be like living with the applicant every day. Is he or she too timid to ask questions? Is he or she a know-it-all? Is he or she honest? Does he or she have poor personal hygiene? There are attributes for which you can't dismiss a candidate (race, religion, national origin, disability, etc.), but there are many personal traits that can create problems at a later date. And, of course, undesirable traits at one company may be perceived as very valuable by a different company.

Forced Reductions

At one time or another, almost every company will find itself in a position where a "workforce reduction" is necessary. Generally, either the company can no longer support the number of employees it has, or a specific worker isn't performing properly or his or her behavior is unacceptable.

It should be noted that most companies have a probation period for new employees of three to six months. During this time, the employee can be dismissed without cause. This trial period allows the employer to determine whether or not a new employee fits in with the rest of the company. If there are personality clashes or if the work product isn't as expected, the employee is quietly dismissed with no blemish on his or her work record.

Laying Off

The most common reason for workforce reduction is financial. The company may see a downturn in revenues, it may outsource jobs, it may be sold, or it may go into bankruptcy.

Assuming that there will be some workers left, you should take care in choosing who to let go and who to keep. The choice isn't always obvious. Some companies, especially those with unions, base layoffs on seniority (even if the engineers are not in a union). Those with the longest tenure are the last to go. This approach generally keeps the oldest and most experienced workers. While this may appear to be the fairest approach, it is clearly arbitrary in terms of value to the company.

Some will choose to keep the workers who are most capable. Likewise, this is not always the best choice, since these workers will generally have an easier time finding a new job. They may not like the new situation and leave.

This means that you may find yourself in the position of needing to hire someone after the department has been decimated. Bringing a new person into such a situation can be very difficult.

The workers that you want to keep are those who will stay and who are capable of performing adequately. In this case, you will want workers that are dependent, conservative, and not self-assured. As a group, these people will find it hard to look for a new job. They may be very capable, but their personalities cause them to want to hang on to what they have, making them less likely to leave. The downside is that they will require more direction and oversight than those employees who are being laid off.

If there is a key person that you really need to keep, you may have to compromise. If that person can easily find another job elsewhere, you must entice him to stay, possibly through a raise or promotion. A promotion is probably a better choice, and appears reasonable as part of a reorganization plan. A contract for a number of years is another possibility, but remember that contracts work both ways. You are obligated to pay his or her salary for that time period.

Termination for Cause

If, after the trial period, someone is incapable of perform-

ing their job or behaves in an unacceptable manner, they may be fired or terminated for cause. This is very traumatic for both the employee and the supervisor, and is akin to a divorce. The worker may have spent years working with the department and now he's being dumped. It's not like Donald Trump pointing a finger at an apprentice and saying, "You're fired."

There are laws in place to prevent arbitrary dismissals. If done improperly, firing someone can lead to lawsuits that can drag on for years. Termination for cause is rarely impulsive or spontaneous (violence or theft are exceptions). You can't fire someone because you don't like him, because his child picks on your child, because he goes to KKK meetings in his spare time, or because he has bizarre sexual habits. You can only fire him because of behavior or performance issues on the job.

Because documentation is necessary, it takes at least six months — and sometimes a year — to get someone out the door. You must demonstrate that actions were taken to make the employee aware of the problem and give him or her a reasonable chance to improve his or her performance. You then need to show that there was no significant improvement over that period of time.

Documentation entails written reports of the problem and written reports of meetings with a third party present as a witness. It's necessary to specify exactly what the problem is to demonstrate that what you are asking isn't impossible. Provide data that shows that everyone else can complete a similar task in certain period of time, and that the problem employee can't. Offer training or other help to improve performance. Provide guidelines and benchmarks for improvement. You will have to have regular meetings with the employee to discuss the situation.

Sometimes, simply asking the problem person to leave will work. However, problem workers often recognize that they will have difficulties securing a new job, so they tenaciously hold on to what they have. They also fail to recognize that being fired from a job, instead of quitting, makes getting a new job much harder. Having "Terminated for Cause" on your job record is a big hurdle to overcome.

Conclusion

When hiring, it's paramount to find the proper fit with the company. By developing a sound job description and using a filtering process that preserves the best candidates, you have a good chance at finding an employee who will fit the company's needs. If you must terminate someone's employment, do it with deliberate care and consideration. **NV**

RoboNexus

The Largest Robotics Event in the Western Hemisphere

October 6-9, 2005
San Jose Convention Center, San Jose, CA

The International Business Development, Educational and Consumer Event for Personal, Service and Mobile Robotics

- 50 Robotics Visionaries and Thought Leaders
- 50,000 Square Foot Exposition
- 5 Comprehensive Professional Development Conferences:
  - Business Development and Opportunities Conference
  - Emerging Robotics Technologies and Applications Conference
  - Robotics Design, Development and Standards Conference
  - Robotics Education and Instruction Conference
  - Consumer Robotics and Entertainment Event

New for RoboNexus 2005!

- Business-to-Business and Consumer Entertainment Expo Areas
- Service Robotics Summit
- Service Robotics Pavilion
- "Robotics at Home" Demo Stage
- Robotics Innovators Awards —the 'Robi'

www.robonexus.com

IDG WORLD EXPO
roboticstrends

Founding Sponsor
iRobot

EXPO Magazine's best new show
2005 Winner

Premier Sponsor
INNOVA HOLDINGS, INC.
Bridging Innovative Technologies

Premier Media Sponsor
POPULAR science

Gold Sponsors
HITEC
THE TECHNOLOGY COLLABORATIVE
MOBILEROBOTS
Phidgets USA.com

Corporate Sponsors
ANIMATICS
CoroWare
evolution robotics
ESS ELECTRONIC SCHOOL SUPPLY, INC.
LEGO education
mBus machineBus.com
Openware Robotics
OLogic
INSPIRED ENERGY
PARALLAX
Robots and Design RND www.rnd.re.kr
RoboTech CENTER
ROGUE robotics www.roguerobotics.com
SEGWAY
ShopBot
SYSTRONIX Embedded Java Spoken Here
SILICON SENSING GyroSensor CRS-03

Media Sponsors
BATTERY POWER PRODUCTS & TECHNOLOGY
CIRCUIT CELLAR
e-DRIVE
Entertainment Engineering
INDUSTRIAL+UTILITY VEHICLE
InfoWorld GET TECHNOLOGY RIGHT
GLOBAL FUTURES
Industrial Robot An International Journal
iToolbox
LINUX JOURNAL
Nuts & Volts
NetworkWorld
PC WORLD
Popular Mechanics
ROBOTBOOKS
sensors
SERVO
VisionSystems DESIGN
TUX

Association Sponsors
Acceleration Studies Foundation
GETI Global Emerging Technology Institute
OMG

Tech Forum

QUESTIONS

I am in need of the following parts for an old, no longer made, Ungarmatic Soldering Station:

Ungar Part #71 - Handle & Cord Assembly
Ungar Part #77 - 700 Deg Farenheit Heater
Ungar Part #79 - Tip Retaining Rings
Soldering Tips - #84, #88, #93

#10051   **Bruce Bubello**
                         Wayne, NJ
                Bruce.Bubello@gmail.com

RadioShack (and others) sells a wireless indoor/outdoor thermometer (#63-1033) that transmits the outside temperature to an indoor numerical display unit. Basically, this is a wireless telemetering system. I would like to convert this unit to telemeter the voltage on an electric fence rather than the outdoor temperature. I have opened the units but there are no markings on the semiconductor units. Has anyone tried such a conversion? Does anyone know how the circuits operate?

#10052   **Emanuel Kramer**
                       Charles City, VA

My HD DVR cable box does not have HDMI — just DVI, but it has SPDIF and Toslink outputs for Dolby 5.1 sound. My digital TV has HDMI and DVI inputs but only analog audio inputs to use with the DVI video input. If I go direct from the cable box to my Dolby Digital receiver, I bypass the audio delay in the TV that compensates for the video processing time delay. The delay is nearly a half second. Does anyone have a circuit for creating a delay for the audio stream? The analog component video of DVD and digital DVI of cable have different delays, so I will need two different delays. I need both optical and SPDIF I/O versions because the TV is optical out only and my DVD is coax only.

#10053   **Dennis Green**
                     Farmington Hills, MI
                     dlgreen@tir.com

I am looking for a circuit that will sum three analog signals to one output. The three inputs are DC, vary from 0 to 1.2 volts, and vary at 0 to 100 Hz. I believe I should isolate the inputs more than with just a simple resistor network. Is there a three-input mixer or amplifier IC that would do this? I'm also open to a discrete-component solution, and would appreciate any suggestions.

#10054   **Frank Helsell**
                         via internet

ANSWERS

[#06051 - June 2005]
I need a transformer that meets the following requirements:
1. Small physical size (toroid/E-core/other)
2. Low/reasonable cost
3. 30:1 Step-up transformer
Input 6 VDC (chopped square wave) running at around 17.5 kHz
Output approx. 180 VDC (chopped square wave)
Input — Primary approx. 15 turns of #24 magnet wire
Output — Secondary approx. 450 turns of #34 magnet wire
If I wind a toroid, what type of material do I use (powered iron core)? I would like to have the size as close to a FT-50 or FT-82 as possible.

Congratulations on basically getting to the solution that you need! The FT-82 core will do the job for the number of turns you propose if the 77-mix ferrite is selected, that is, an FT-82-77 core. Now, the longer version of the answer.

Most of the information required to come up with a definitive answer is contained in an almost "D" size document divided into 24 pages from Amidon Associates, "Iron-Powder and Ferrite Core Forms," dated April 1989. I was able to find some, but not all, of this information also on the **www.amidoncorp.com** website.

Ferrite or iron powder core? Ferrite supports a higher flux density, resulting in a smaller core and lesser number of primary turns. Ferrites work better than iron powder cores at the lower frequencies — lower than tens of MHz. The 15 primary turns that you have chosen will work well with the larger size ferrite core which you have chosen.

Will 15 turns work? The formula below was extracted from page 9 of the above document; though, I could not find a website link. This formula gives the flux density produced in a core in gauss, given the variables listed below. A_e, the cross section area of the magnetic path, comes from printed page 14, or the link listed in the second table below. It varies with the size of the core.

The recommended flux density that we can subject the core to is a function of frequency. High flux density can overheat the wire and the ferrite core material. Page 21 lists flux density (gauss) vs frequency: 3,000 gauss at 10 kHz, 2,000 gauss at 20 kHz, 1,200 gauss at 40 kHz. We could pick 2,000 gauss at 20 kHz, since it is close to your 17.5 kHz. We can get a slightly better number by interpolating 17.5 kHz between 10 and 20 kHz:

$(20 -17.5)/(20-10) = (2000-x)/(2000-3000)$
$2.5/10=(2000-x)/(-1000)$
$-(2.5/10)1000 = 2000-x$
$-400-2000 =-x$
$x=2400$

We do not want to exceed 2,400 gauss in the equation below. If you do not like this number, fall back to the more conservative 2000 gauss.

$$B_{max} = \frac{E_{PK} \times 100}{4.44 \ A_e \ N \ f}$$

E_{PK} = peak applied RMS voltage
A_e = core cross-section area in cm^2
N = number of turns
f = frequency in MHz

core	A_e	N	E_{PK}	f	B_{max}
FT-50-77	0.133	15	6.0	0.0175	3800
FT-82-77	0.246	15	6.0	0.0175	2000
PC-1408-77	0.25	15	6.0	0.0175	2060
EA-77-188	0.225	15	6.0	0.0175	2280

Below the B_{max} equation we list the parameters for your too small FT-50-77 core. With B_{max} = 3,800 gauss from the equation, we exceed the capability of this core (2,400 gauss). We plug in the values for the larger FT-82-

77 core and get an acceptable 2,000 gauss (anything less than 2,400 gauss).

Will the #34 and #24 windings fit? Though I found website data for pot-cores and E-cores, for the lower table, I found none for the toroids. Though, printed page 4 lists 219 turns per layer of a similar T80 core. Your 450 turns of #34 wire will require a little over two layers to wind, easily fitting the 0.52" opening. Fifteen turns of #24 wire will occupy a fraction of the possible 66 turns per layer.

To summarize, your choice of 450 turns of #34 and 15 turns of #24 wire on an FT-82-77 core looks like a winner. However, the length of 450 turns of hair fine wire will be a challenge to thread through the toroid hole. You might consider using a pot-core or E-core, where you can wind the coil on a bobbin, around which the pot-core or E-core is assembled.

The B_{max} table (shown first) lists a PC-1408-77 pot-core as having an acceptable 2,060 gauss flux density by evaluating the equation. We also list an EA-77-188 at a passable 2,280 gauss. Moving to the "turns" table, we list (from the website) the number of turns of #34 or #24 wire which fits the bobbin.

core	turns #34	turns #24	www.amidoncorp.com/
FT-82-77			see datasheet
PC-1408-77	275	25	aai_potcores.htm
PC-1811-77	425	50	
PC-2213-77	700	100	
EA-77-188	669	79	aai_ecores.htm
EA-77-250	1240	147	

The PC-1408 will not accommodate 450 turns of #34 wire, nor will the next larger size, PC-1811. The still larger PC-2213 will accept 700 turns #34 or 79 turns #24. We propose to use:

$450/700 + 15/100 = 0.67 + 0.15 = 0.72$

That is 72% of the available winding space. This will fit if machine wound or carefully layer wound with 28% space left to spare. If scramble wound, we need a 25% margin, which we have.

Let's look at fitting the windings into an E-core. We calculate a flux density of 2280 gauss for the EA-77-188 (first table). The EA-77-188 will accept 669 turns of #34, 79 turns of #24. We will use:

$450/699 + 15/79 = 0.64 + 0.19 = 0.833$

The windings will fit if machine wound. If random wound by hand, we need 25% margin, but only have 17%. Use the next larger E-core, EA-77-250, to be sure.

We have not included some of the larger core sizes in the B_{max} calculations because the larger A_e values for bigger cores will produce smaller (passing) flux density numbers.

Where to go from here? Any of the larger core sizes for

which we have not calculated B_{max} values could be recalculated with fewer than 15 primary turns of wire to see if we can reduce the size of the windings. The "D" size data sheet, which I refer to, used to be shipped with orders for cores from Amidon Associates. This document might clear up any further questions you may have. In fact, you need it to follow this letter.

Dennis Crunkilton
Abilene, TX

[#09052 - September 2005]

I need a power supply circuit that will take a voltage and invert it. The input voltage can range from 0 to 3 volts. The output will range from 0 to -3 volts and supply up to 3 amps of current.

#1 First, to be humorous, if you think you can get 3 amps at 0 volts, you might have overlooked simply reversing the polarity of the wires.

Second, I get positive and negative 3 volts from a 3V solar panel using an ICL7660 IC, which takes 3 volts and gives -3 volts at 100 milliamps, which charges two out of the four cells that are capable of supplying 3 amps at 5 volts.

Third, there are new chips and old germanium transistors that can work at less than 1.5 volts, and be used as oscillators for voltage doublers.

William Como
Bethpage, NY

#2 Figure 1 is a current limited, 0 to -3V mirror circuit. The positive supply is necessary for the op-amp to control the LM333K down to 0 volts, and to

allow a current dump of at least 5 mA when the LM333K is delivering 0 volts of output. The capacitors are tantalum, and C1 is not necessary if the LM333K is less than 4 inches from the power supply filter capacitor. The accuracy is very dependent on how well R1 and R2 are matched values.

Chuck Larson
Largo, FL

Figure 1

Current limited voltage mirror for zero to negative 3 volts out, 3 amps.
Negative supply = 5.7 to -15 volts.

Don't let the Exos get the Axo, *sign up today!*

TETSUJIN 2005

Tetsujin 2005 — originally scheduled for October 6-9 at RoboNexus — has been postponed. The registration deadline passed on June 13th, and there were not enough entries to hold the competition on the scheduled date.

The event will be rescheduled for 2006. Individuals and teams interested in competing should drop us an email to get on the competitors list for updates and info regarding the competition. The entry deadline will remain open for now. Don't send money or the entry form yet, just email to let us know you're interested and we'll keep you posted on what to do next and when to do it.

Hey, look on the bright side, not only do you *STILL* have time to enter, you have even *MORE* time to design and build your exosuit! So, there's no excuse now NOT to compete! This is a hard competition and certainly one worth participating in. If it were easy, everyone would be doing it. This is your chance to stand apart from the crowd and flex your engineering and bot-building muscles. So, hook up with your friends, fellow engineers, classmates, geeksquad, gearheads, whoever, and form a team. No need to go it alone!

If you haven't already done so, check out the new rule set and event changes and start scribbling out your preliminary design. There's no time to lose! Send an email to **tetsujin@servomagazine.com** with your name and email address and we'll make sure you get all the latest info and news.

CHALLENGE 1:

Weightlifting. Ascend stairs in your suit to the lifting platform and lift a load of from 100 to 1,000 lbs* from a squatting position to a height of at least 24 inches*, return the load to the ground in a controlled manner, and descend the stairs. Stair-climbing may be unpowered. The winner is the competitor who lifts the most weight.

CHALLENGE 2:

Dexterity. Stack nine concrete cylinders weighing about 70 pounds each in a 4-3-2 vertical arrangement, but don't knock them over as the pyramid grows! The winner is the competitor who arranges the cylinders in the shortest time.

CHALLENGE 3:

Walking Race. Walk the 100 foot* long U-shaped challenge course, stepping over a small obstacle at the half-way point. The shortest time wins, with a time bonus being granted based on any auxillary load carried. Walking must be powered.

The current rule set is available online at

www.servomagazine.com

Questions can be directed to

tetsujin@servomagazine.com

Don't wait, sign up for one, two, or all three challenges!

Specifics of the competition are in a tentative state and may be subject to change.

SX ToolKit PRO

SX Toolkit PRO; #45182; $229.00

With the Professional Development Board and the SX-Key®, your project will know no limits! The SX Toolkit PRO is well-equipped with an SX-Key Manual v2.0, Daubach's Programming the SX Microcontroller book, an SX-Key programmer, 2 SX28 DIP chips, 2 resonators, programming cable and CD-ROM. The SX Tech Toolkit PRO is the best starting point to get you developing SX projects with Parallax's SX-Key programming tool.

The SX-Key® is Parallax's primary development tool for the SX line of microcontrollers from Ubicom, supporting every SX chip. The SX-Key programming tool supports in-system programming and source-level debugging. The new SX/B compiler lets BASIC Stamp® users have quick success with the SX.

The Parallax Professional Development Board (PDB) is high-quality, fully-integrated development platform for BASIC Stamp®, Javelin Stamp, and Ubicom SX28AC/DP microcontrollers. A wide variety of typical I/O (LEDs, LCD interface, buttons, etc.) devices and circuitry are built into the PDB, providing the developer with an ideal platform for rapid microcontroller project development.

Order online at www.Parallax.com/SX
Or call toll-free 888-512-1024 (Mon-Fri, 7am-5pm, PT)

PARALLAX

BASIC Stamp and SX-Key are registered trademarks of Parallax, Inc. Parallax and the Parallax logo are trademarks of Parallax, Inc.